

DTIC FILE COPY

IMPACT AND POTENTIAL OF DECISION RESEARCH ON DECISION MAKING

REPORT OF A DEPARTMENT OF DEFENSE RESEARCH ROUNDTABLE WORKSHOP

AD-A132 276

Submitted by:

Office of Naval Research
and
Office of Naval Technology

Under Order No. N00014-87-G-0146

DTIC
SELECTED
NOV 17 1987
D

Submitted by:

Approved for public release
Distribution Unlimited

American Psychological Association
1200 17th Street, N.W.
Washington, D.C. 20036

and

Decision Science Consortium, Inc.
1695 Preston White Drive
Suite 300
Reston, Virginia 22091

December 1987

This work relates to Department of Navy Grant N00014-87-G-0146
issued by the Office of Naval Research. The United States
Government has a royalty-free license throughout the world in
all copyrightable material contained herein.

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION Unclassified			1b. RESTRICTIVE MARKINGS		
2a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release; distribution unlimited		
2b. DECLASSIFICATION / DOWNGRADING SCHEDULE					
4. PERFORMING ORGANIZATION REPORT NUMBER(S)			5. MONITORING ORGANIZATION REPORT NUMBER(S)		
6a. NAME OF PERFORMING ORGANIZATION American Psychological Association		6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION Department of the Navy, Office of Naval Research		
6c. ADDRESS (City, State, and ZIP Code) 1200 17th Street, N.W. Washington, D.C. 20036			7b. ADDRESS (City, State, and ZIP Code) 800 North Quincy Street Arlington, Virginia 22217		
8a. NAME OF FUNDING / SPONSORING ORGANIZATION Off. Nav. Research & Off. Nav. Technology		8b. OFFICE SYMBOL (If applicable) 1142PS	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER Grant # N00014-87-G-0146 "		
8c. ADDRESS (City, State, and ZIP Code) 800 North Quincy Street Arlington, Virginia 22217			10. SOURCE OF FUNDING NUMBERS		
			PROGRAM ELEMENT NO 61153N 62234N	PROJECT NO RS34H20 & RT4424804	TASK NO
					WORK UNIT ACCESSION NO
11. TITLE (Include Security Classification) Impact and Potential of Decision Research on Decision Aiding					
12. PERSONAL AUTHOR(S) Martin A. Tolcott, Ph.D., and Virginia E. Holt					
13a. TYPE OF REPORT Final Technical Report		13b. TIME COVERED FROM 3/15/87 TO 3/14/88		14. DATE OF REPORT (Year, Month, Day) 1987, December	
15. PAGE COUNT 65					
16. SUPPLEMENTARY NOTATION Prepared in cooperation with the Decision Science Consortium, Inc.					
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)		
FIELD	GROUP	SUB-GROUP	decision research, decision aids, human judgment, decision making, cognitive abilities, behavioral research, cognitive tasks, decision performance, decision aided technology		
19. ABSTRACT (Continue on reverse if necessary and identify by block number) A workshop was held on May 4-5, 1987 to explore the past impact and future potential of decision research on decision aid design, and to recommend a research agenda that would contribute to decision aid effectiveness. Major conclusions were: 1. Research on human judgment and decision making has produced rich findings and theories about cognitive abilities and fallibilities, but designers of decision aids often fail to take advantage of this knowledge. 2. Decision aid research and development has been dominated by the push of currently popular technologies rather than by systematic study of user needs, resulting in widespread lack of confidence and frequent rejection of sophisticated aids by potential users. 3. New research is needed on agenda setting, option generation in unstructured situations, pattern recognition in decision making, effective use of graphic displays, achieving user trust, and avoiding restricting decision aid applications to narrowly focused problems without requiring large knowledge bases. 4. Managerial efforts are needed to achieve better cooperation among researchers, decision aid developers, and potential users.					
20. DISTRIBUTION / AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT <input type="checkbox"/> DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION Unclassified		
22a. NAME OF RESPONSIBLE INDIVIDUAL Gerald S. Malecki, Office of Naval Research			22b. TELEPHONE (Include Area Code) 202/696-4741		22c. OFFICE SYMBOL Code 1142PS

TABLE OF CONTENTS

	Page
PREFACE	3
SUMMARY	5
1.0 INTRODUCTION	7
2.0 RESEARCH ON HUMAN COGNITION	8
3.0 RESEARCH ON DECISION AIDS	10
4.0 MANAGERIAL ISSUES	12
SCHEDULE	13
PARTICIPANTS	15
PRESENTATION SUMMARIES	
Amos Tversky	17
Robyn Dawes	19
Baruch Fischhoff	23
Randall Davis	27
Dennis Leedom	31
Marvin Cohen	35
Herbert Simon	39
Stephen Andriole	53
Jeffrey Grossman and Capt. J.R. Fitzgerald	55
D. Warner North	57
Karl Weick	61



Accession For	
NTIS CR&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution	
Availability Codes	
Avail and for	Special
A-1	

PREFACE

This document is the report of a Workshop on the Impact and Potential of Decision Research on Decision Aiding, the second in a series of Department of Defense Research Roundtables. The Workshop was sponsored by the Office of Naval Research (ONR) and the Office of Naval Technology (ONT), organized and conducted by the American Psychological Association (APA) and Decision Science Consortium, Inc. (DSC), and held at the National Academy of Sciences on May 4-5, 1987.

The idea for the Research Roundtable series was generated by the APA Committee on Research Support (CORS), a committee reporting to APA's major scientific policy-making body, the Board of Scientific Affairs. CORS initiated the series to provide a forum for the discussion of behavioral research issues related to the defense mission. Decision research and decision aiding was chosen as the second topic for consideration. The First Roundtable, "Issues in Psychological Research and Application in Transfer of Training," held on February 27-28, 1986, was organized by the American Psychological Association in cooperation with the Federation of Behavioral, Psychological, and Cognitive Sciences, and sponsored by the Army Research Institute for the Behavioral and Social Sciences.

CORS set the framework for the Roundtable project, and appointed members Gary M. Olson and Baruch Fischhoff to provide oversight for this workshop. Subsequently, a planning committee developed specific guidelines and selected presenters. Planning committee members were Gary M. Olson (University of Michigan), Baruch Fischhoff (Carnegie Mellon University), Gerald S. Malecki (ONR), Stanley C. Collyer (ONT), Virginia E. Holt (APA), and Martin A. Tolcott (DSC).

The individuals selected as presenters reflected a variety of interests and experience in the areas of decision research and decision aiding. They represented the disciplines of behavioral decision research, artificial intelligence and expert systems, statistical decision theory and organizational theory, and included researchers from academic settings as well as those who have developed and used decision aids or decision-aiding procedures in military and non-military settings. In addition, the planning committee invited a number of individuals from government laboratories involved in decision-aiding research and development, to actively participate in the discussions. The workshop program and a complete listing of workshop participants, including those who made presentations, may be found on pages 18-22.

Presenters from the research community were asked not only to describe relevant research but to specify how it has been or could be used to improve decision performance. Those from the development community were asked not only to describe decision-aiding techniques but to specify the kinds of research findings they had found useful. All presenters were asked to identify promising future research areas.

Presenters were asked to provide a summary of their remarks in advance for distribution to all the participants. These summaries are included in this document. One presenter, Herbert Simon of Carnegie-Mellon University, did not prepare a summary; in its place this document includes the report to the National Research Council of the Research Briefing Panel on Decision Making and Problem Solving (1986), which he chaired. The summaries will be found starting on page 23.

This report was prepared by Decision Science Consortium, Inc. It is based on the summaries, on notes taken during the presentations and ensuing discussions, and on comments submitted by the presenters after reviewing a draft version. It is organized according to three major themes that seemed to recur: research on human cognition, research on decision-aiding technology, and managerial techniques for strengthening the links between them.

Virginia E. Holt
American Psychological Association

Martin A. Tolcott
Decision Science Consortium, Inc.

SUMMARY

A workshop was held on May 4-5, 1987 to explore the past impact and future potential of decision research on decision aid design, and to recommend a research agenda that would contribute to decision aid effectiveness. Major conclusions were as follows:

1. Research on human judgment and decision making has produced rich findings and theories about cognitive abilities and fallibilities, but designers of decision aids often fail to take advantage of this knowledge. Although the relevance of laboratory research to complex real-world decisions may be questioned, procedures to minimize bias in human judgment are applicable in many situations, and their incorporation into a computer based tactical decision aid has been demonstrated. There is mixed evidence about the extent to which laboratory findings can be generalized to trained personnel working in their domain of expertise, and research is needed to clarify this issue. In addition, there is little indication that the direction of behavioral research has been guided by decision-aiding needs.

2. Decision aid research and development has been dominated by the push of currently popular technologies rather than by systematic study of user needs, resulting in widespread lack of confidence and frequent rejection of sophisticated aids by potential users. Operationally successful aids have tended to be simple tools for handling relatively structured, but onerous or repetitive, tasks. Decision aid applications would be facilitated by a commonly accepted taxonomy of cognitive tasks matched to analytical techniques, and decision aid research should be directed toward gradual progress in moving from structured to more high-level, relatively unstructured decision tasks, and toward aids for use in organizational settings. In addition, improved methods are needed for testing or validating decision aid effectiveness, and indeed for measuring the quality of decision performance.

3. New research thrusts are needed in the relatively ignored areas of:

- agenda setting, or the focusing of scarce attention and the priority ranking of information and decisions in a high work-load environment;
- option generation in unstructured situations, possibly viewed as the creation of new knowledge from old;
- pattern recognition in decision making, especially the effective training of pattern recognition skills;
- more effective use of graphic displays in eliciting knowledge structures and helping develop decision strategies;
- methods of achieving user trust in decision aids;
- methods to avoid restricting decision aid applications to narrowly focused problems without requiring impossibly large knowledge bases.

4. Managerial efforts are needed to achieve better communication and cooperation among researchers, decision aid developers, and potential users. These should include:

- systematic review of the current state-of-the-art and applications of decision aiding technology;
- a directed exploration of the scope for actual and potential research to enhance this technology;
- a coordinated program of research and development oriented toward aids that would best meet high priority operational needs.

5. Finally, it was intended that the workshop address not only research on human cognition and decision aiding, but research on normative models (logical, mathematical and statistical) that contribute to decision aiding technology. However, these latter issues were largely ignored. This may suggest their relatively low priority, but it may also be a function of unwitting bias in the selection of presenters and participants. Any inference that normative research was discussed and regarded as unimportant, would not be warranted.

1.0 INTRODUCTION

There is a growing interest in the development of techniques to improve decision making—if not necessarily the quality of decisions, at least the process by which decisions are made. Decision aiding takes a variety of forms: one-on-one relationships between decision analysts and clients, computational aids in individual decision-maker workplaces, decision-support systems designed to serve many decision makers in an organization, and expert systems designed to capture human knowledge in a given problem domain and ultimately to replace the human completely. The domains or contexts in which decision aiding is being investigated have expanded significantly, from early interest in one-time management decisions and recurring medical diagnoses in well-specified domains, to public policy decisions involving public attitudes toward risk, and to military situation assessment and command/control decisions in submarines, aircraft, ships, shore-based command centers and Army battlefields. This broad potential for utilization is testimony to the increasing recognition of the importance of cognition, planning, and problem-solving in a wide variety of situations, and the danger of serious performance degradation or failure due to information overload and rapid-response requirements.

The past fifteen years have also witnessed significant progress in research on human decision making. This research reflects several different areas of interest and approaches, and has resulted in important advances in both descriptive and prescriptive theories. Behavioral decision research has tended to emphasize descriptive formulations of individual inference and choice, and powerful new theories (e.g., prospect theory, framing, heuristic reasoning) have emerged to explain how decision makers process information and why they typically demonstrate certain errors and biases (by which is meant deviations from so-called normative models). Mathematical/statistical approaches have produced new or improved prescriptive models and logical paradigms for handling inference and choice in the absence of comprehensive or reliable human assessments, and have paved the way for development of artificial intelligence and expert systems approaches to decision aiding. Organizational theorists have characterized how decisions are made in actual organizational settings and the institutional factors that impede effective decision making and constrain the means for improving it.

To what extent has the design of decision aids benefited from the kinds of research described above? Have some research findings already been used effectively in techniques that have been shown to improve decision making? Are there additional findings that could be incorporated into decision aids, with high potential for some kind of positive effect? Could decision aids be significantly improved by certain lines of research that have not yet been undertaken? This workshop was organized and conducted to explore these questions.

Although the formal presentations were necessarily given sequentially and emphasized different aspects of these issues, three major themes tended to recur in both the papers and the ensuing discussions. The first was the nature of research on human judgment, cognition and decision behavior, its value as a basis for designing decision aids, the failure to exploit existing research findings, and the need to enhance our knowledge in certain areas and to initiate research in new areas. The second was the tendency of decision aid development to be based on the push of technology, especially current fads, rather than on systematic study of user task requirements or behavioral characteristics, the inadequacy of techniques for evaluating decision aids, and the resulting reluctance (at least in the military) to move decision aids from research and development to operational use. The third theme was the nature of the administrative or managerial roadblocks to achieving a more effective interaction between the research and the decision-aid development communities.

Because of the predominance of these themes, this report has been organized around them rather than in accordance with the sequence in which the presentations were given. Reference to individual presentations will be made as appropriate, and the individual presentation summaries (contributed by the authors) are included on pages 23-77.

2.0 RESEARCH ON HUMAN COGNITION

Several of the speakers emphasized the importance of research on human cognitive abilities and fallibilities as a basis for determining requirements for decision aiding. However, the ensuing discussions revealed that only rarely have the findings from this research been incorporated into decision aids, even those that are strictly in the R&D phase, and much less so into operational systems. Much of the discussion was concerned with whether the research was irrelevant or the development community did not understand its significance; arguments were made for both points of view.

Tversky illustrated the nature of human judgmental fallibility by reference to two phenomena that have been found to be quite robust:

- (1) framing effects—the finding that logically equivalent choice problems produce very different responses when worded (or “framed”) as gains (e.g., survival rates) as compared with losses (e.g., mortality rates);
- 2) elicitation effects—the finding that preferences are not simply revealed but actually constructed during the elicitation process, and are in fact influenced strongly by the elicitation method used.

Although early research on these topics was conducted in university laboratories with untrained subjects and relatively artificial (or unfamiliar) problems, there is some evidence that even trained respondents may exhibit these trends when the content of the problem is in their area of expertise. Although the evidence here is mixed, what can be concluded about the role of experience is that there is no guarantee that experience will remove the effects.

In most realistic situations the problems are more complex than those used in this research; choices and preferences are often implicit rather than explicit, and biased judgments may be difficult to identify. When explicit judgments are called for, the aiding technique recommended by Tversky is to frame the problem and elicit the judgments in several different ways, allowing for a sensitivity analysis to both measure the effect and bring it to the awareness of the decision maker. In at least two situations, the relevance of this research to decision aiding is obvious: one is in the elicitation of subjective probabilities and utilities during a formal decision analysis; the other is the elicitation of probabilistic knowledge from experts for development of an expert system or AI knowledge base. In both cases, judgments are explicit, minimization of bias is critical, and a sensitivity analysis based on multiple framing and elicitation approaches could result in significant improvement. Greater effort is needed to identify appropriate applications of this research in other situations where the judgments are implicit and the relevance not so obvious.

Dawes reaffirmed the existence of widespread judgmental biases, such as ignoring of base rates and vulnerability to framing effects, adding that even training on the framing effect does not reduce that bias. He emphasized that resistance to the use of decision aids is based on humans' reluctance to acknowledge that their reasoning abilities are flawed. As a technique for increasing people's acceptance of their cognitive limitations as well as helping to overcome them, he recommended wider use of graphics such as Venn diagrams or influence diagrams, which he has found give a better indication of base rates and show more clearly than verbal descriptions the relationships among uncertain factors. He cited experimental evidence showing that the use of influence diagrams in fact persuaded experts to change their judgments about a problem in their own area of expertise.

Fischhoff stressed the need for decision makers to construct models of the problem situation they are facing, and the importance of providing aids for the model-building process itself, to help ensure the use of the most appropriate model and an understanding of its limits. He referred to the applicability of cognitive psychology research and procedures to the model-building process, and urged a strengthened interaction between the cognitive

science and the decision research communities. In this connection the discussion brought out the view that protocol analysis or process tracing—a method widely used by cognitive psychologists to identify the underlying causes of errors in reasoning—is not generally used in decision research, and that its use might help explain the findings.

The issue of whether experts show the same limitations as novices was explored in more depth by Davis, who was concerned with human behavioral problems in the development of expert systems. In particular, he stressed the need for improved methods and theories of knowledge elicitation that would help ensure that the knowledge obtained was not subject to experts' biases, such as might result from overconfidence, efforts to appear rational, inarticulateness in the face of overwhelmingly large knowledge bases, etc. He called for a better theory of uncertainty, one that would accommodate non-numerical representations of degree of belief and allow them to be integrated with numerical combination mechanisms. The discussion also brought out the need to find ways of extending the applicability of AI to relatively ill-structured decisions.

Davis' presentation introduced the topic of group decision behavior by suggesting that multiple experts may provide a better knowledge base than single experts, but he cautioned that we lack techniques for knowing when that is the case, for eliciting knowledge from groups, and for reconciling diverse judgments.

Weick dealt more extensively with group decision making in an organizational context. He emphasized the importance of pre-decisional activities aimed at stabilizing the (usually turbulent) organizational environment and enhancing the group's understanding of the problem to be solved. The relevant cause and effect relationships can be clarified by the construction of a "cause map," similar in concept to the Venn or influence diagram mentioned earlier.

Perhaps the most controversial idea advanced in the workshop was one of Weick's suggestions for learning about the problem; namely, that direct action should be taken early and used more frequently than it usually is, since action generates both environmental stability and information. Other recommended techniques were: thinking about the problem as if it were in the past, ensuring diversity in the group composition, reduction of organizational stress, and embedding the problem in a broader context. Discussion brought out the need to validate these ideas in realistic settings, a theme that was raised frequently in connection with decision aids generally. Also, the need for increased research attention to group decision making, especially when the group members were distributed, was repeatedly mentioned.

Although all the presenters dealt with research needs to some extent, Simon focused directly on this issue, dealing with human behavioral research relevant to both individual and group decision making and decision aiding. His overall theme was that understanding decision making requires a broad look at the whole process rather than simply that of choice among options. Conventional decision aiding techniques, such as operations research, subjective expected utility, decision trees, etc., may be useful in well-defined and relatively quantifiable situations. However, research has largely ignored the ill-structured parts of the process. He stressed two aspects that especially need research attention: agenda setting (deciding what decisions must be made), and the generation of options.

Simon described the agenda-setting issue in terms of the setting of priorities and the allocation of attention (a scarce resource) to the relevant information, as exemplified by the sifting and ranking of information in military intelligence analysis. Research in this area should investigate human attention and information-filtering processes. The research should center on the user, and the information systems should be designed for an environment in which attention is scarce and non-numerical information, such as natural language, must be accommodated.

With regard to option generation, Simon exemplified the problem by reference to higher levels in an organization dealing with unstructured problems, and described it as the creative process of putting together new sets of alternatives (new knowledge) from existing building blocks. Research is needed to gain a better understanding of the process, to develop AI systems that can generate alternatives, and to develop training techniques that will promote creativity in humans.

Another research area identified by Simon was that of problem representation; he reiterated the potential benefit of graphics in helping people not only understand relationships, but actually perceive new information that

points toward a solution. On a related issue, he stated that the recognition of patterns is often essential in military decision making, but that training in pattern recognition is poor and could be benefited by appropriate research.

In connection with distributed decision making, Simon felt that key research issues were the extent to which components or nodes can be separated, and the most effective way to allocate functions.

The presentation by Cohen provided a balance between the emphasis on human behavior and that on decision-aid research. He described an approach to the design of decision aids that is based on: (1) research findings about the limitations of human cognition; (2) investigation of user mental models about the specific problem area being addressed; and (3) incorporation in the aid of features that are both personalized in that they accommodate user-preferred information-processing styles, and prescriptive in that they guide the user away from error-producing procedures and biased results. The key is to provide displays that represent information compatibly with the user's own internal representation, while a normative model, running in parallel, provides a prompt whenever its results differ significantly from the user's or the user is adopting a procedure that could lead to error.

Since the method relies heavily on cognitive research findings, Cohen's recommended research agenda emphasized such research and is therefore included in this section. He calls for research on knowledge structures and decision making strategies used in various problem domains and their linkage to cognitive biases; methods for reducing biases that are compatible with user-preferred strategies; techniques for knowledge elicitation that facilitate mapping of knowledge onto displays; and richer normative models relevant to the constructive and creative processes used by real decision makers.

3.0 RESEARCH ON DECISION AIDS

Aside from the general sense that decision aid designers often ignore the results of research on human behavior, there were several recurrent themes bearing on improvements in decision aiding. It was repeatedly pointed out that research and development for computer-based decision aids have been dominated by the push of current technology rather than the pull of explicit user requirements; that aids have generally failed to engender user trust; that many aids attempt to address high-level, relatively sophisticated decision problems while there are many unsatisfied real needs (in the military, at least) at lower levels where tasks are onerous and time-consuming; and that aids are rarely, if ever, carried through to systematic evaluation. Recommendations emerging from these presentations and discussions centered on applied rather than basic research issues, and in some cases dealt more with administrative or managerial procedures (the latter will be covered in Section 4.0).

Leedom's presentation emphasized the need to design decision aids (especially expert systems) that are compatible with the user's point of view, and called attention to the inadequacy of current methods of eliciting and representing expert knowledge to achieve this goal. AI system design is currently driven by such things as programming convenience, familiarity with specific techniques and personal preferences of the developers, and little attention is given to cognitive consistency with the user. Achieving cognitive consistency is an especially difficult problem in an organizational or distributed context because group members must develop a common perception of the problem and a common set of semantic constructs in the face of limited communication channels. The use of only one or two experts to construct a knowledge base is insufficient because of potential differences among them, and even the identification of experts worth "cloning" is an unresolved issue.

Leedom pointed out the complex and ill-structured nature of most real decision tasks, and the failure of current decision aids to solve such problems. He stressed the need for users to understand the aiding system and have grounds to trust it, and deplored the fact that decision aid development usually stops with a demonstration rather than a systematic validation, in large part because validation methods are inadequate.

Discussion brought out the sense that the greatest operational success has been achieved with relatively low-level computational aids that relieve humans of routine, well-structured calculations, and that there are many unmet needs of this type that seem to be of little interest to the decision aid R&D community. The requirement to develop a taxonomy of battle management tasks to which a variety of aiding technologies can be matched, was identified as a prerequisite to developing more successful operational aids.

Andriole developed in more detail the ideas that decision-aid development is vulnerable to the push of currently popular technologies, and that a taxonomy to permit more effective matching of tasks to analytical methods is required. He pointed out that over the past few years there has been a trend toward shifting of task control from humans to computers, and that this trend has been roughly paralleled by a shift in the popularity of various aiding technologies, from decision analysis, to operations research, to AI, and (currently) to biological emulation (neural networks). Several recent examples of task misallocation and force-fitting of technologies were given.

There have been many proposed taxonomies of human tasks, but few that deal in detail with cognitive tasks, and little attention has been devoted to showing relationships between such tasks and analytical techniques that could aid them. Andriole presented several tentative taxonomies—of decision tasks, of analytical methods, and of possible matches between them—to illustrate his ideas, but recommended a research program to investigate the strengths and weaknesses of various matchings in greater depth to provide a more reliable basis for design decisions. He reaffirmed the importance of basing system design on user requirements, but advocated early and continued user involvement in the design process (rapid prototyping and evaluation) rather than a formal, up-front requirements analysis, as more appropriate for modern complex systems.

Grossman's presentation also emphasized the importance of satisfying user requirements. His co-author, Captain J.R. Fitzgerald, Commander, Destroyer Squadron 31, who was to have given specific examples of aids used in the context of anti-submarine warfare operations and their strengths and weaknesses, was unable to attend. Nevertheless, Grossman was able to represent the views of the operational personnel who have used those aids. He reaffirmed the importance of user confidence, and the tendency for users to prefer simple aids that have wide applicability, rather than sophisticated (automated) aids that tend to be highly context-dependent and specialized in the problems they can handle.

He suggested that, to enhance user acceptance, decision aids should be: synergistic (result in better and faster decisions), robust (resistant to information loss and show graceful degradation), reliable (consistent), adaptable (work under variable situations), responsive (capable of self-assessment), flexible (permit easy updating), user-friendly (easy to learn and operate), and testable (permit testing at all levels of realism and degradation).

Grossman identified three other dimensions along which one might discriminate between decision aids that tend to be operationally successful and those that have gained the attention of the R&D community: the currently successful aids have been those that: (1) operate in a fixed (vs. variable) decision-making environment; (2) use centralized or organic (vs. decentralized or distributed) sources of data; and (3) supply a deterministic (vs. probabilistic) output. He suggested that R&D efforts should move out from the successful toward the challenging end of these dimensions, rather than focusing on the extreme challenges. Finally, he pointed out that a single decision aid for two team members working together is a different entity than one for each of them, and recommended research on how such aids can best be designed for an organizational setting.

North represented the discipline and practice of formal decision analysis, which is based on normative theories and models of subjective expected utility (SEU) maximization; it relies heavily on the elicitation of probability and utility estimates from the decision maker, and makes extensive use of such graphics as influence diagrams and decision trees. The effective use of decision analysis, according to North, usually requires the presence of a decision analyst, and many computer-based models for the construction and analysis of decision trees are now available commercially as aids to the analyst.

Such models may be generic in nature (i.e., essentially templates that can be applied to a variety of decisions) or may be developed in the course of aiding a decision maker on a specific problem. North provided examples of both types, illustrating applications to foreign policy (Persian Gulf), and domestic policy (synthetic fuels commercialization and power plant emission control). Components of decision analysis such as multiattribute utility (MAU) analysis and decision tree diagrams have been investigated in connection with military decision aids, but the number and difficulty of the judgments involved have precluded their acceptance for most military applications.

Discussion brought out the close relationship between the judgments called for in decision analysis and those required in developing expert systems. Decision analysts' techniques for highlighting uncertainty and asking questions in several ways (to minimize cognitive bias) were stressed as potentially valuable in work on expert systems. In fact, it was suggested that a decision-analytic model generated in the course of solving a specific problem might be regarded as an expert system (or at least as the repository of required expert judgments) in that problem domain. This point of view might facilitate research that could lead to the desirable objective of making decision analysis more independent of experienced analysts.

Finally, discussion brought out again the deficiency in efforts to evaluate decision aids and the inadequacy of methodology for doing so; concern was expressed for the difficulty of solving a prior problem, namely, the evaluation of complex decisions themselves in the face of their heavy dependence on underlying subjective judgments. Clearly, evaluation criteria beyond the satisfaction of the decision maker were felt to be needed.

4.0 MANAGERIAL ISSUES

Aside from specific research recommendations discussed in Sections 2.0 and 3.0, two general themes recurred: 1) existing basic research knowledge is not being effectively used, and 2) much R&D in this area is not addressing real user needs. The discussions did not necessarily focus on managerial steps that might be taken to address these problems, since that was not the objective of the workshop. Nevertheless, from time to time, suggestions emerged (or more properly, were implied) during the formal discussions and to some extent during informal conversations between sessions. In this section an attempt is made to describe briefly some of the ideas that were mentioned.

For decision aid developers to make more effective use of existing basic research knowledge requires that they be informed of the research findings in a way that makes apparent the relevance of existing knowledge to the problems they are addressing, the potential benefits that might accrue from applying this knowledge, and an estimate of the costs (or time, or effort) involved in achieving implementation. Some evidence of support from the user community would be helpful, even at this relatively early stage. How is this informing of the developing agencies to be achieved?

Traditionally the R&D community relies primarily on distribution of written technical reports, and secondarily on formal briefings (at technical meetings and program reviews), to exchange this kind of information. The technical reports are written by the researchers themselves who, with some exceptions, are interested mainly in the scientific or theoretical aspects of their work, and are usually not in the best position to understand its potential applications, much less to convince others. Technical briefings and program reviews typically have crowded agendas that do not allow time for extensive discussion.

Most effective communication among busy people is achieved by face-to-face discussion. Suggestions were made at the workshop that the researchers and developers must be brought into closer touch with each other. Periodic visits to each other, and longer-term exchange programs, are not uncommon and are often productive; they should be expanded. But it would appear that better use could also be made of the "middlemen" in the government agencies, those individuals who both sponsor the research and are themselves in close touch with development efforts (and sometimes with the users). These individuals are probably in the best position to bridge the gap between research and development, and should be given time and sufficient travel funds to increase their direct and informal contacts with both groups so that they could more effectively serve as the communication links in both

**WORKSHOP ON THE IMPACT AND POTENTIAL OF DECISION RESEARCH
ON DECISION AIDING
MAY 4-5, 1987**

Schedule

Monday, May 4

- | | |
|------------------|---|
| 8:00 am | Registration and Continental Breakfast |
| 8:30 | Welcome and Introduction
Virginia E. Holt, American Psychological Association
Gerald Malecki, Office of Naval Research
Martin A. Tolcott, Decision Science Consortium, Inc. |
| 8:45 | "Framing the Problem and Eliciting Beliefs"
Amos Tversky, Stanford University |
| 9:45 | "Rationality in Deciding How to Decide"
Robyn Dawes, Carnegie-Mellon University |
| 10:45 | Coffee Break |
| 11:00 | "Building Models with Decision Aids"
Baruch Fischhoff, Carnegie-Mellon University |
| 12:15pm | Lunch |
| 1:15 | "Knowledge-based Systems as Decision Aids: What Have We
Got? What Do We Need?"
Randall Davis, Massachusetts Institute of Technology |
| 2:15 | "Issues in Knowledge Representation"
Dennis Leedom, Army Research Institute |
| 3:15 | Coffee Break |
| 3:45 | "When the Worst Case is Best: Mental Models, Uncertainty, and Decision Aids"
Marvin Cohen, Decision Science Consortium, Inc. |
| 4:45-5:15 | General Discussion |
| 6:30 | Group Dinner |

Tuesday, May 5

8:00 am	Continental Breakfast
8:30	"Research Opportunities" Herbert Simon, Carnegie-Mellon University
9:30	"Optimal Human-Computer Task Allocation" Stephen Andriole, George Mason University
10:30	Coffee Break
10:45	"Decision Aids - Who Needs 'Em?" Jeffrey Grossman, Naval Personnel Research and Development Center
12:00 pm	Lunch
12:45	"Decision Aids for Decision Makers: Views of a Decision Analysis Practitioner" Warner North, Decision Focus, Inc.
1:45	"Interpretation-Based Decision Aids for Organizations" Karl Weick, University of Texas
2:45	General Discussion
3:15	Adjourn

PARTICIPANTS
Workshop on the Impact and Potential of Decision Research
on Decision Aiding

May 4-5, 1987
Washington, D.C.

Stephen Andriole
Department of Information Systems
and Systems Engineering
George Mason University
Fairfax, VA 22030

Rex Brown
Decision Science Consortium, Inc.
1895 Preston White Drive, Suite 300
Reston, VA 22091

Marvin Cohen
Decision Science Consortium, Inc.
1895 Preston White Drive, Suite 300
Reston, VA 22091

Randall Davis
School of Business
Massachusetts Institute of Technology
Cambridge, MA 02139

Michael Drillings
Army Research Institute
Office of Basic Research
5001 Eisenhower Avenue
Alexandria, VA 22333-5600

Baruch Fischhoff
Department of Social and Decision Sciences
Carnegie-Mellon University
Pittsburgh, PA 15213

Jeffrey Grossman
Navy Personnel Research and
Development Center
San Diego, CA 92125

R. F. Hudson
Naval Air Development Ctr.
NADC Code 3C
Warminster, PA 18974-5420

Donald E. Bamber
Navy Personnel Research &
Development Center
Code 41
San Diego, CA 92152-6800

Susan Chipman
Program Manager
Cognitive Science, 1142CS
Office of Naval Research
Arlington, VA 22217-5000

Stanley C. Collyer
Office of Naval Technology
Code 222
800 N. Quincy Street
Arlington, VA 22217-5000

Robyn Dawes
Department of Social and
Decision Sciences
Carnegie-Mellon University
Pittsburgh, PA 15213

Jon J. Fallesen
Army Research Institute Field Unit
P.O. Box H (PERI-SL)
Fort Leavenworth, KS 66027-0347

Robert A. Fleming
Navy Personnel Research and
Development Center (HFESG)
1411 South Fern Street
Arlington, VA 22202-2896

Virginia E. Holt
American Psychological Association
1200 17th Street, N.W.
Washington, D.C. 20036

Clinton W. Kelly
DARPA/ISTO
1400 Wilson Boulevard
Arlington, VA 22209-2308

Dennis Leedom
Army Research Institute
for the Behavioral and Social Sciences
5001 Eisenhower Avenue
Alexandria, VA 22233

Alan Meyrowitz
Office of Naval Research
Code 1133
800 N. Quincy Street
Arlington, VA 22217

Warner North
Decision Focus, Inc.
4894 El Camino
Los Altos, CA 94022

Judith Orasanu
Army Research Institute
Basic Research Office
5001 Eisenhower Avenue
Alexandria, VA 22333

Herbert A. Simon
5818 Northumberland
Pittsburgh, PA 15217

Yale Smith
Rome Air Development Center
Decision Aids Branch
Griffiss AFB, NY 13441-5700

Martin Tolcott
Decision Science Consortium, Inc.
1895 Preston White Drive, Suite 300
Reston, VA 22091

John R. Valusek
Air Force Institute of Technology
AFIT/ENS
Wright-Patterson AFB, OH 45433

Steven F. Zornetzer
Office of Naval Research
Code 114
800 N. Quincy Street
Arlington, VA 22217

Gerald S. Malecki
Office of Naval Research
Perceptual Science Program
Code 1142PS
800 N. Quincy Street
Arlington, VA 22217-5000

Judy N. Moracco
NAVSEA-CEL-MP
Naval Sea Systems Command
CEL-MP 27
Washington, D.C. 20362

Gary Olson
Department of Psychology
University of Michigan
Ann Arbor, MI 48109 Ann

Sharon L. Riedel
ARI Field Unit
P.O. Box H (PERI-SL)
Ft. Leavenworth, KS
66027-0347

James G. Smith
Office of Naval Research
Code 1241
800 N. Quincy Street
Arlington, VA 22217-5000

Michael G. Sovereign
Joint Command, Control &
Communications Curriculum
Code 74
Naval Postgraduate School
Monterey, CA 93943

Amos Tversky
Department of Psychology
Stanford University
Stanford, CA 94305

Karl E. Weick
Department of Management
CBA 4.202
University of Texas
Austin, TX 78712

Framing the Problem and Eliciting Beliefs

Amos Tversky
Department of Psychology
Stanford University

There is a general agreement among students of decision making, including decision theorists, management scientists and psychologists, that people often need help in making difficult choices and that the quality of human decision making could be enhanced by the development of an appropriate support system. Many students of decision making also share the hope that the recent development of both hardware and software computer technology could provide the basis for such systems. In order to develop effective decision aids, however, we need to understand, at least in general terms, what are the major factors that limit human decision and judgment and introduce error and bias. At present, we have only partial answers to this question, and we cannot exclude the possibility that the answers will turn out to be very specific so that different support systems may be needed for different tasks. Some general information regarding human limitations have emerged from the growing field of behavioral decision research over the last two decades. The following paragraphs summarize two general phenomena that appear relevant to the construction and design of decision aids.

1 *Framing Effects* Recent research has demonstrated that alternative framing of the same options could lead to drastically different choices (see Bazerman, 1986; Kahneman & Tversky, 1984; Tversky & Kahneman, 1986). Moreover, these effects are large, systematic and prevalent; they are observed in the decisions of both naive and sophisticated respondents, with both monetary and nonmonetary outcomes. For example, the proportion of experienced physicians who chose radiation therapy rather than surgery as a treatment for lung cancer rose from 16% to 50% when the problem was framed in terms of mortality rates rather than survival rates. Recent studies of loss aversion indicate comparable effects when the same outcomes are framed as gains or as losses.

2 *Elicitation Effects* In contrast to the traditional analysis of choice that treats preferences as given, research suggests that people do not generally have ready-made values and that preferences and beliefs are constructed—not merely revealed—in the elicitation process. Moreover, because preferences (like attitudes) are constructed from vague impressions, different constructions can give rise to different responses. Thus, choice is contingent, or context-sensitive, in the sense that different methods of elicitation give rise to different decisions. In particular, preferences can be inferred from direct choice between options or from a matching procedure (e.g., pricing) in which the decision maker adjusts one option to match another. Recent investigations (Tversky, Sattath & Slovic, 1987) show that the more important dimension looms larger in choice than in matching. This discrepancy produces preference reversals (Slovic & Lichtenstein, 1983) as well as other nonoptimal patterns of choice.

The failures of invariance induced by framing and elicitation effects underscore the need for decision aids. At the same time, they emphasize the contingent character of human decision making, which poses a serious challenge to existing decision aids. The difficulties revealed by the psychological analyses of judgment and choice, I believe, cannot be readily resolved by extending the decision maker's memory or computational capabilities. They may require a different approach to decision analysis that recognizes the lability of choice and the problem of resolving conflict.

Bibliography

- Bazerman, M.H. (1986). Judgment in managerial decision making. John Wiley & Sons, Inc.
- Kahneman, D., Slovic, P., & Tversky, A. (Eds.). (1982). Judgment under uncertainty: Heuristics and biases. New York: Cambridge University Press.
- Kahneman, D., and Tversky, A. (1984). Choices, values and frames. American Psychologist, 39, 341-350.
- Tversky, A., & Kahneman, D. (1986). Rational choice and the framing of decisions. The Journal of Business, 59, Part 2, S251-S278.
- Slovic, P., & Lichtenstein, S. (1983). Preference reversals: A broader perspective. American Economic Review, 73, 596-605.
- Tversky, A., Satath, S. & Slovic, P. (1987). Choice, matching and the reversal of preferences. Unpublished manuscript, Stanford University, Stanford, CA; Hebrew University, Jerusalem, Israel; Decision Research, Eugene, OR.

Rationality in Deciding How to Decide

Robert M. Dawes
Department of Social and Decision Sciences
Carnegie-Mellon University

Probability theory as we now understand it was first developed in the Italian Renaissance and later formalized and extended in France in the seventeenth century. Focused toward a frank concern on deciding which gambles are favorable or unfavorable, it developed normative principles for determining probability (which originally referred to wise courses of action not derivable from "first principles") in structured situations of uncertainty. It consequently developed as a form of mathematics. Then, in 1944, von Neumann and Morgenstern published the Theory of Games and Economic Behavior,¹ which again focused on decisions, generalizing the concept of a gamble to that of an alternative course of action with probabilistic consequences. They demonstrated that if certain intuitively compelling axioms about choice between such alternatives were satisfied, then a decision based on expected utility was mandated, where the measurement of utility was determined by the choices themselves.

Historically, much subsequent work was based on a distinction between rationality and expected value as normative criteria for decision making, as opposed to descriptive theory of how people and organizations actually behave. Various "paradoxes" (e.g., Elsberg's) were discussed by theorists who were concerned whether the von Neumann and Morgenstern system constituted a normatively compelling one (and the independence and comparability axioms were challenged), while theorists concerned with the descriptive application of the system asked whether people and organizations actually behave according to its tenets. The empirical research provided a resounding answer to this descriptive question: No.

Specifically, people often use "cognitive heuristics" in reaching decisions under uncertainty; such heuristics include but are not limited to: satisficing searches, elimination by aspects, probability estimation on the basis of availability of instances (either evaluated by the number that can be "brought to mind" or by the difficulty of doing so), conclusions based on anchoring and (under) adjustment, beliefs in probability and plausibility of explanations on the basis of the ease with which they can be constructed in the imagination and so on. Most of these heuristics constitute a form of "bounded rationality," because they work better than nothing, but can lead to predictable and systematic biases in many situations. All violate the expected utility principle, and many violate principles of rationality itself; for example, representativeness leads to equating inverse probabilities in the absence of a corresponding belief in equal base rates; framing effects lead to contradictory choices in what people recognize and admit to be the same decision making situation, and so on.

While there are very strong trends to violate normative principles according to the structure of these heuristics, however, such principles are not always violated. Consider, for example, the following passage from Clark's Biography of Bertrand Russell (1976).² Russell's grandmother, in an effort to dissuade him from marrying his first wife, impressed upon him how much insanity there was in his family. Nine years later, he was considering having children and consulted a doctor about the heredity component of insanity. Clark (pages 96-97) writes:

Four days later he saw his doctor, who said "it was my duty to run the risk of conception, the fear of heredity being grossly exaggerated. He had said 50% of the insane have alcoholic parentage, only 15% insane parentage. This seemed to settle the matter." Settle, that is, until Russell, the potential parent, was overtaken by Russell the statistician; the footnote in his journal reads: "They didn't say what proportion of the total population are insane and drunken respectively, so that his argument is formally worthless."

We are not all as spontaneously smart as Bertrand Russell was. But can we be? My point is that we do not always think one way when making decisions. While the distinction between normative and descriptive has some value, it is not absolute. We often accept normative principles if they "come to mind" and their relevance to a particular decision making situation is understood. That statement is descriptively true, and it is sometimes. Occasionally, we

accept the normative principles without questioning, despite a counter intuition; for example, if I multiply an odd by an even number in my head and come up with an odd number, I know I made a mistake. Sometimes, we accept the normative principles ambivalently; for example, suppose that I have 100 applicants for a position as my secretary and I would be happy with any one of the top 5. I can work out mathematically that if I randomly searched through only half these applicants, my probability of coming across one of these is slightly over .97. (Whether my evaluation system will allow me to recognize such a person is another matter). I may ambivalently truncate the search, or I may not. (All my calculations necessarily lead me to conclude is that I need not search through 96 in order to find one of the top 5 with a virtual certainty). Sometimes, we inconsistently adopt normative principles and reject them. For example, when I was President of the Oregon Psychological Association, some of my colleagues wished to change the laws waiving confidentiality in order to report suspected child abuse—which mandated such reporting if the child might be in any danger—to include reporting of any suspicion of child abuse whenever or wherever it might have occurred: “Because one thing we know about child abusers is they never stop on their own.” “How do you know that?” “Because I’ve dealt with over 50 child abusers, and none of them had stopped on their own.” “How did you see them?” “They were mandated by the court.” “If they did it just once, what is the probability they would end up in court?” “Very, very low.” After this questioning my clinical colleagues acknowledged that their sample did not provide a basis for their conclusions. But then the next month, their experience with the abusers they saw clinically being so compelling, they would propose the same change.

Most notably, we can decide how to decide—and our method of decision can be evaluated a priori on the degree to which it follows normative principles. Here is where I hope the future of decision aiding lies. It can help us to decide in a manner that is rational (not necessarily involving acceptance of the criterion of maximizing subjective expected utility). It can, at the least, point out the systematic departures from rationality (as well as from that criterion). Insofar as decision aids can assist the decision maker to reach conclusions in accord with normative principles, and accept these principles, these aids could have tremendous benefits.

(I do not claim expertise in the field of constructing decision aids [“expert systems”] meant merely to “simulate” whatever the decision maker does—or however it is that people think. I do, however, believe that there is a basic problem with such an approach. It assumes that decision and thought follow single principles, but we do not decide or think in just one way. Russell, for example, understood the importance of base rates long before the field of “Bayesian” decision making emerged. Others didn’t. Nor is it true that a “superior” person will always make good decisions,³ and particularly not somebody whose alleged superiority is based on “expertise,” which is essentially a socially defined variable. Thus, even attempts to simulate the decision and thinking processes of a single individual picked as unusually good will not necessarily aid anyone, even that individual).

I predict that the main opposition to decision aids will be based on our firm belief that our intellects per se are not flawed. We are perfectly willing to recognize that there are limits on our perceptual ability (e.g., that we cannot see microbes—and hence understand bubonic plague—without a microscope), or our physical abilities, and even on our motivational abilities to overcome such problems as laziness or “weakness of the will.” But it is difficult to accept the idea that our logical reasoning abilities per se are so limited that they need to be aided by some external device in the same manner that a doctor reaches a diagnosis on the basis of external report from a blood laboratory. (For example, shortly after the first atomic weapons were used, many people hypothesized that the human race would destroy itself because it had become “too smart,” rather than “too stupid” to control its own technology; and even now most military leaders seek maximal power and response flexibility without concern that they are intellectually incapable of coping with that very same power and flexibility.) In addition, there is the tradition ranging from Plato through the Catholic Church to Freud that mistakes and biases are to be ascribed purely to the “base motivations” interfering with an otherwise god-like intellect. Accepting decision aids presupposes an acceptance of our basic cognitive limitations, which will involve a radical change of belief about ourselves.

¹von Neumann, J. and Morgenstern, O. (1944). Theory of Games and Economic Behavior. New York: Wiley.

²Clark, R.W. (1976). The Life of Bertrand Russell. New York: Knopf, pgs. 96-97.

When [Governor, Supreme Court Judge, well-known "liberal"] Earl Warren was testifying in favor of interning Japanese Americans before a Congressional Committee on February 21, 1942, a questioner pointed out that there had been no instances of sabotage or other types of espionage by Japanese Americans up to that time: Warren's response: "I take the view that this lack is the most ominous sign in our whole situation. It convinces me more than perhaps any other factor that the sabotage we are to get, the Fifth Column activities we are to get, are timed just like Pearl Harbor was timed. I believe we are being lulled into a false sense of security."

Building Models with Decision Aids

Baruch Fischhoff

Department of Social and Decision Sciences
Carnegie-Mellon University

Decision theory offers a highly flexible language for modeling decision making problems. If one can speak the language, then one can mobilize the theory's resources for solving those problems. One longstanding research program within behavioral decision theory has involved assessing and, where needed and possible, improving decision makers' fluency. Its focus has been on decision analysis, the approach to decision aiding most explicitly concerned with utilizing judgments provided by decision makers and their aides. However, the general methodology, and even some of its substantive results, are applicable to other decision aiding techniques that rely on flexible modeling languages. Such languages might include cost-benefit analysis, probabilistic risk analysis, and "spread-sheet."

The computerization of these techniques offers both possibilities and pitfalls. The added computational power increases the size of the problems that can be tackled. The development of networks facilitates pooling the judgments of diverse and distributed experts. The standardization of procedures allows routine use of best methods for eliciting or expressing information. On the other hand, computers can put the techniques in the hands of the "masses," absent the steadying hand of practiced analysts who know their limitations. Complexity can mean analyses that defy evaluation, both before and after decisions are made. Standardization may be done by computer specialists, insensitive to the cognitive processes and limitations of users.

Like natural languages, these modeling languages are better suited to capturing some phenomena than others. As a result, using them requires an understanding of their limitations, including an ability to read between the lines for things left unsaid, or said imprecisely. Apropos the language analogy, much of the research in this area has adopted the methods of cognitive psychology, examining how people acquire or exercise the various intellectual skills needed to use a language.

A large share of this research has concerned one particular skill, the ability to summarize the extent of one's knowledge in a quantitative expression of confidence (or uncertainty). This concentration reflects the essential role of uncertainty in decision making, the distinctive "subjectivist" perspective of decision analysis, and the existence of opportunities for validating such judgments. These studies have produced fairly robust, and fairly simple, patterns of results. Current research here is focused on ways to improve performance through iterative elicitation procedures, through systematic provision of feedback, through engineering the evaluation of evidence, and through training for substantive expertise (i.e., will making people more knowledgeable also help them to know the limitations of their own knowledge).

An area of growing interest has been the elicitation of values, or the tradeoffs that decision makers wish to make between competing objectives. A natural default assumption is that people know what they want. As a result, much research has focused on developing simple ways of asking people about their values and sophisticated ways of expressing the subtleties of those values in formal ways. More recently, there has been growing concern about an accumulating number of "quirks," cases in which people's expressed values are affected by seemingly irrelevant features of how questions are asked. Some of these quirks reflect predictable "framing" effects, whereas others reflect nuances of the associations evoked by the context within which questions are posed. They crop up in behavioral economics and survey research, as well as in decision making research. They seem to represent cases in which decision makers may have strong opinions, but not necessarily for the precise question that they are required to answer. In such cases, they may be unduly sensitive to some aspects of question formulation and unduly insensitive to others. They pose a methodological and philosophical challenge to decision aiding, which is getting some uneasy attention. Particularly troubling are difficulties in determining what is the right question to ask and fears that the elicitation procedure will shape (and not just capture) decision makers' values.

The basic concept of decision theory is that of an individual decision maker. However, decision aids are often used in organizational settings, in which the knowledge and concerns of varied individuals must be collected, integrated, and distributed. In such settings, it is essential to ensure that the participants speak the same language, not only that of decision theory, but also that of particular problems. Without face-to-face interactions, inconsistent usage may go undetected, leading to confused actions. The problems of deriving and implementing consensual definitions are also receiving attention.

Once the terms of a decision have been set, the model itself must be communicated, including some understanding of the limiting assumptions of the pieces that have been supplied by different sources within the organization. Moreover, that shared understanding must be maintained over time, as the model is refined and updated with inputs from different quarters. Work is just beginning into how the limitations of individual decision makers are ameliorated or exacerbated by embedding them in organizational settings and connecting them with decision aids.

Decision theory is a formal theory. As a result, it has little to say directly about the substantive side of problems like defining terms or evoking the right considerations with a value question—beyond requiring clarity and consistency. Nor is it helpful for determining what action options should be considered in the first place. Perhaps predictably, research into these aspects of the decision-making (and model-building) process have lagged. Work on option generation and formulation are particularly important for decision problems where the “space” of options is ill-structured. Help is needed here.

Other hard, but important topics include: how to express the limits of models, how to handle off-model considerations, how to help decision makers match archetypal models to specific situations that they face, and how to integrate decision theory and its models with expert systems.

Bibliography

A. Three essays that expand on some of these ideas from my own perspective are:

Fischhoff, B. (1987). "Judgmental aspects of risk analysis." in V. Covello & M. Merkhofer (eds.), Handbook of risk analysis, prepared by the National Science Foundation for the Office of Management and Budget.

Fischhoff, B., and Cox, A. (1985). Conceptual framework for regulatory benefits assessment. In J. Bentkover, V. Covello, and J. Mumpower (Eds.), Benefits Assessment: The State of the Art. New York: D. Reidel.

Fischhoff, B., and Johnson, A. (1987). "The possibility of distributed decision making. In J. March and R. Weissinger-Baylon (Eds.), Political-military decision making. New York: IAI Press.

B. Three recent handbook chapters give somewhat different overviews of current research in his field:

Levi, A. and Abelson, R. (1984). In Handbook of social psychology. Reading, MA; Addison-Wesley.

Slovic, P., Lichtenstein, S., and Fischhoff, B. (in press). In Handbook of experimental psychology. New York: Wiley.

Fischhoff, B., Svenson, O., and Slovic, P. (1987). In Handbook of experimental psychology. New York: Wiley.

C. Three book-length treatments on the relationship between decision theory and behavioral research are:

von Winterfeldt, D., and Edwards, W. (1986). Decision analysis and behavioral research. New York: Cambridge University Press.

Watson, S., and Buede, D. (1987). Decision Synthesis. Cambridge, UK: Cambridge University Press.

Fischhoff, B., Lichtenstein, S., Slovic, P., Derby, S., and Keeney, R. (1981). Acceptable risk. New York: Cambridge University Press.

D. Issues of probability elicitation are reviewed by:

Wallsten, T., and Budescu, D. (1983). Encoding subjective probabilities: A psychological and psychometric review. Management Science, 29: 151-173.

E. Problems of value elicitation are discussed in:

National Research Council. (1982). Survey measure of subjective phenomena. Washington, D.C.: Author.

Tversky, A., and Kahneman, D. (1981). The framing of decisions and the rationality of choice. Science, 221: 453-458.

All of these publications provide references to the primary empirical literature.

Knowledge-Based Systems as Decision Aids: What Have We Got? What Do We Need?

Randall Davis
Massachusetts Institute of Technology

Knowledge-based systems (KBS) have been a focus of considerable interest recently for their utility in aiding (and making) decisions in a variety of fields. I will speak as a builder of such systems, approaching the issue from two basic directions: What have we got to offer? What do we need? In considering what we have to offer, I will review very briefly knowledge-based systems technology as it currently stands, noting the places where it has drawn from previous decision aiding research and the places where it has differed. The intent here is to make clear what ideas and technology are currently available for building decision aid systems using this approach. In considering what we need, I'll suggest places where our existing stock of ideas is most pressingly incomplete and speculate on the kinds of research that would help.

What Have We to Offer?

A First-Order Theory of Expertise

Work on GPS (the General Problem Solver) was motivated by what might be called the O-th order theory: the belief that the distinguishing characteristic of intelligence was its generality and that expertise could be captured by a single (or a few), domain-independent methods like means-ends analysis. Work that began in the mid 1960's soon demonstrated that power in problem solving arose far more from task specific knowledge; simply put, experts are experts because of what they know. This has led to the current focus on knowledge representation and accumulation in AI, cognitive science, and other areas.

Ways of Building Non-numeric Models of Thought Processes

Models of all varieties allow us to express and test what we think we understand. Knowledge-based systems (and AI programming techniques generally) facilitate the construction of models of thought that are non-numeric and focused on process rather than end product. One important result of building such models is that they are runnable, and hence readily testable, offering the potential for a significant increase in mental hygiene.

The View of KBS as Repositories of Expertise

A knowledge-based system is more than just a program that happens to produce good answers. It is as well a repository for our current understanding of the problem, much like a survey article or detailed monograph. Unlike a book, however, it is particularly malleable and particularly dynamic, much simplifying the task of constant growth and updating of that understanding.

Some Simple Mechanisms for Knowledge Representation and Reasoning

The current stock of ideas includes rules, frames, logic, and other schemes, along with a variety of corresponding reasoning techniques. We also have some broad guidelines about when to use each. Working systems form an existence proof of the adequacy of these tools on a narrow range of tasks.

Some Simple Ideas on Knowledge Acquisition

Work has been done on debriefing experts, both manually and with various levels of automation, resulting in the development of some basic but effective techniques. Our current stock of tools are suggested primarily by enlightened common sense, leading to such techniques as "knowledge acquisition in context": present the expert with a run of the program whose outcome he disagrees with, show him the system's line of reasoning that lead to that result, and ask him to specify where the reasoning went wrong.

A Very Basic Approach to Uncertainty

Faced with the difficulty of the task, KBS have for the most part taken a pragmatic approach to the problem of inexact reasoning, using methods of defining and combining certainty that are intuitive and simple and that provide tolerable performance. Once again working systems offer an existence proof of the adequacy (though certainly not optimality) of these tools.

What We Need

We Need to Know When the Experts are Lying

Well, at least when they're likely to be exaggerating. Sitting at the heart of much of our work are the intermingled beliefs that experts are experts because they know something, and that we can capture that knowledge and put it to work in our programs. This has led to the widespread practice of debriefing experts, yet numerous papers suggest that experts routinely report to be more confident of their answers than the outcomes suggest they should be. How can we deal with this apparently pervasive bias? More broadly, how do we know an expert is really an expert, particularly on tasks like medical diagnosis, where definitive answers are not always available?

We Need More Powerful Knowledge Acquisition Techniques

Enlightened common sense has given us some good ideas, but surely we can develop more sophisticated techniques. The problem arises in two forms: articulate and inarticulate expertise. Articulate expertise is that which the expert can express; the problem here is one of volume—how can we help the expert navigate through the vast body of what he knows, so we can record it; how can we help him be methodical and remember it all? A second set of issues arises in relying on the expert's description of his own reasoning process. The tendency is to make the process sound more rational and logical in the re-telling than it was in the original. How can we get more accurate descriptions of process? Inarticulate expertise—what the expert knows but cannot describe—presents a different set of interesting problems. Techniques like multidimensional scaling, hierarchical clustering, and the ID3 family of induction programs may prove useful here, but all are subject to the weakness that they are purely syntactic, leading them at times to suggest concepts and rules that have no meaningful interpretation.

We Need a Normative Theory of Multi-Expert Decision Making

The suggestion that knowledge-based systems are built by debriefing a single expert typically elicits the response that committees are more effective, so why not debrief multiple experts and have them interact. The answer is that we don't know how. Nor is it obvious that multiple experts are necessarily better than one. Under what circumstances is that true? How can we make it true, i.e., how *should* experts interact? The HEARSAY speech understanding system offered one interesting model of how to get multiple experts to cooperate. What other models might there be, that are similarly founded on a particular insight about effective problem solving?

We Need a Better Theory of Expertise

Perhaps the basic premise of our first order theory is wrong. Perhaps experts really are different from the rest of us, in ways that are not accounted for by what they know. Dreyfus, for example, argues for a five stage model of expertise, with claims that some levels are inherently unreachable by mechanical information processing systems. His argument is not compelling, but it is important that we consider such alternatives to our basic assumptions. Alternatively, what if the premise is correct and knowledge really is "all there is." The view is becoming pervasive: a recent overview article on expertise in a cognitive science volume talks about "...the necessity to focus on the organization and structure of knowledge in psychological research." Must we all become epistemologists eventually?

We Need a Better Theory of Uncertainty

We need one that is *transparent* and *operational*. Transparency refers to the ability of a system to make its line of reasoning clear, typically done in knowledge-based systems by reviewing the sequence of inferences made. This works to some degree, but can easily be defeated if the uncertainty handling mechanism is opaque and obscure.

We also need to consider theories of uncertainty that are not limited to using numbers to represent what we know. While numeric theories (like Bayes) have the advantage of a simple and uniform combination mechanism, they overlook the possibility that there may be different kinds of reasons to believe or disbelieve something, with corresponding different ways of combining them. While numeric theories can be used to capture the effect of such insights, they do not motivate us to think in such terms. Simply put, numbers may be too simple a representation for something as potentially complex and multifaceted as uncertainty.

We Need to Know How to Make Robust Decisions

Historically, the emphasis has been on optimal decisions. Yet the dynamic nature of almost any real-world problem makes this an ephemeral hope: what good is an optimal decision if the problem changes before you have a chance to carry it out. What good, for example, is an optimal set of gate assignments in an airport in the face of real world contingencies like mechanical breakdowns, traffic delays, weather delays, crew changes, etc.? Optimal answers are often brittle, in the sense that small changes to the problem result in sharp decreases in their utility.

Can we instead consider ways of arriving at decisions that are *robust*, in the sense that they are very good (though not necessarily optimal), yet relatively insensitive to small changes in the problem? Sensitivity tests can help us to recognize such results, but how can we generate them?

Bibliography

Buchanan. (1986). Expert systems in the literature. Expert Systems, 3, (1), 32-51.

This article reviews a wide range of knowledge-based systems, categorizing them by class of problem solved and calibrating the stage of completeness they have reached. The paper has 366 references, making it a useful starting point for finding out about almost anything going on in the area as of January 1986.

Davis, Buchanan, Shortliffe. (1977, February). Production rules as a representation. Artificial Intelligence, pp. 15-66.

This is a quite detailed and technical article describing the workings of the original Mycin system, one of the first rule-based systems. It deals with a wide range of the program's capabilities, including explanation and some of the basics of knowledge acquisition.

Fikes and Kehler. (1985, September). The role of frame-based representation in reasoning. CACM, pp. 904-920.

A good recent overview of some of the ideas involved in using frames as a knowledge representation.

Hayes-Roth. (1985, September). Rule-based systems. CACM, pp. 921-933.

A recent overview of rules as a knowledge representation.

Genesereth and Ginsberg. (1985, September). Logic Programming. CACM, pp. 933-941.

Moving AI Technology to the Organizational Level

Dennis K. Leedom
U.S. Army Research Institute
for the Behavioral and Social Sciences

The Present Situation

The representation of expert knowledge is an issue that pervades most development efforts within the related fields of decision aiding and artificial intelligence. Over the past several years, researchers have espoused a variety of approaches for eliciting and codifying expertise; yet, they have reached little agreement over exactly which methods are most appropriate. In fact, there even exists little agreement over the criteria one should use in selecting a particular representational approach.

As a result, the choice of paradigm for a particular decision aid is often made solely on the basis of programming convenience, familiarity with a specific technique, or personal preference of the programmer. Little consideration is given to the needs of the actual military user—i.e., achieving “cognitive consistency” with the training, experience and perspectives of the user.

Coincident with this practice is the fact that little attention is given during the development of a specific decision aid to (1) a validation of its knowledge base or (2) an assessment of how that knowledge is used (or misused) by the military user. To date, emphasis has been placed more on the “demonstration” of novel modeling structures or reasoning processes, rather than on an objective evaluation of the decision aid’s performance under actual operational conditions.

The Need for Cognitive Consistency

The need for consistency between the representation of expertise in a computer-based decision aid and the expertise of the military user can be seen at two different levels. At the man-machine level, human decision makers retain ultimate responsibility for their decisions. While this fact is generally true in everyday life, it is specifically emphasized during combat operations. As a result, the user must both understand and trust the decision aid with respect to underlying assumptions and perspectives that influence the aid’s output. Unfortunately, most problems from the battlefield are complex and ill-structured. And, because users are rewarded for solving the *right* problem, emphasis is necessarily placed on developing appropriate constructs and problem frames. Decision aids that offer only one perspective on a problem domain may represent both an elegant and useless tool for the military decision maker.

At the organizational level of military operations we see the need for a second type of consistency. Decision making responsibility is typically distributed across a staff. As a result, staffs spend a great deal of their time developing and maintaining a common perception of the problem domain. A prerequisite to this process is the development and maintenance of a common semantic framework. Quite simply, entities and relationships on the battlefield must mean the same thing to different individuals, if those individuals are to effectively collaborate and exchange information.

This has typically been achieved in two ways: military training and communication. Training, however, is only a partial answer since expertise is only loosely organized and codified at the conceptual level of battle management. During actual combat operations, communication channels among the various decisionmakers are at a premium for the exchange and reconciliation of ideas and information. In the future, however, communication channels will become increasingly scarce as staffs are geographically distributed and communication systems are threatened both physically and electronically. What staffs require for the future are efficient languages and knowledge-based systems for communicating ideas and meaning.

This last idea can be understood more clearly by looking at the role of expertise in battle management. One of the primary mental tasks of a commander's staff is to interpret or make sense out of the environment. The expertise of the staff serves to suggest appropriate constructs and problem frameworks for identifying and exploiting decision opportunities. In the future, this task may be a shared responsibility between man and machine. As a result, this cooperative process must be supported by knowledge-based technology that can serve in both a decision aiding and a communication capacity. That is, decision aids will be used both to develop insight into a specific battle management problem and to efficiently communicate that insight to another staff element.

Design Criteria: 1975 versus 1987

In light of these identified requirements, it becomes necessary to evolve a new set of design criteria for decision aiding technology. In 1975, emphasis was placed primarily upon:

- o Correspondence of the symbolic representations to objects and relationships in the real world
- o Generation of additional facts and truth maintenance
- o Linkage of units and structures within the model
- o Matching of structures for equality and similarity
- o Self-awareness

In 1987, we must concern ourselves with additional criteria in order that what is produced is seen by the military user as both valid and useful:

- o Facilitation of multiple perspectives
- o Consistency with the user's training and experience
- o Organization and coherency of the knowledge
- o Robustness in face of real world ambiguities
- o Efficient articulation and communication of complex ideas

Knowledge Representation Paradigms

A variety of approaches are available for encoding and representing expertise. In terms of symbolic representations, one can approach the issue from a bottom-up, generative perspective or a top-down, interpretative perspective. Generative models have traditionally been organized around the use of first-order predicate calculus (or extended derivatives of this logic formalism). Knowledge is assumed to be organized at an atomic level and is generally considered to be context independent.

Limitations and inefficiencies with purely logical systems have led to more interpretative methods such as semantic networks, frames, scripts, conceptual graphs, and entailment meshes. Such approaches not only provide more efficient organization of specific types of information, but also permit consideration of multiple perspectives and open worlds (i.e., "beliefs" and "assumptions" rather than just "true/false" values). A drawback with these more interpretative approaches is that truth maintenance and consistency are often sacrificed in the process.

With advances in parallel computation machinery, it is possible that paradigms such as parallel distributed processing will offer new approaches for organizing and representing semantic knowledge. At issue here is not only the task of making computers more efficient, but also the task of organizing knowledge bases in a manner more closely resembling that of human memory.

Research Agenda

At least in the foreseeable future, there does not seem to exist any single method suitable for all forms of knowledge representation and decision aiding. As a result, attention must be focused on a careful matching of methods to the specific demands of the problem domain. This involves the careful identification of relevant constructs for battle management and the coherent characterization of battle management decisions.

For decision aiding technology to ultimately be useful, greater attention needs to be placed on the empirical validation of the knowledge bases in an operational setting. We must move beyond "technology demonstrations" to systems that provide "value added" to the staff's decision making process. Elements of this validation process must necessarily attend to the following issues:

- o Multiple perspectives
- o Cognitive consistency with the user
- o Efficient articulation and communication of complex ideas

Finally, it is now possible to begin looking at decision aiding technology as a part of the general evolution of organizational knowledge and culture. As part of this process, attention needs to be focused on the following issues.

- o Expanding the areas of public knowledge and semantic meaning in specific areas of battle management
- o Utilizing knowledge-based systems as a vehicle for institutional memory
- o Evolving decision support systems coincident with staff education and training

Bibliography

- Johnson-Laird, P.N. (1983). Mental Models: Towards a cognitive science of language, inference, and consciousness. New York: Academic Press.
- McClelland, J.L. and Rumelhart, D.E. (1985). Distributed memory and the representation of general and specific information. Journal of Experimental Psychology: General, 114, 159-188.
- Schank, R. and Abelson, R. (1977). Scripts, plans, goals, and understanding. Hillsdale, NJ: Erlbaum.
- Schvandeveltdt, R.W., Anderson, M., Breen, T.J., Cooke, N.M., Goldsmith, T.E., Durso, F.T., DeMaio, J.C. and Tucker, R.G. (1983). Conceptual structures in fighter pilots. AFHRL-TP-83-2, Flying Training Division, Air Force Human Resources Laboratory.
- Sowa, J.F. (1984). Conceptual Structures. Reading Mass.: Addison-Wesley.

When the Worst Case is Best:
Mental Models, Uncertainty, and Decision Aids

Marvin S. Cohen
Decision Science Consortium, Inc.

Consider a pilot faced with uncertainty regarding the presence of an enemy surface-to-air missile installation on his planned flight path; and, for contrast, consider an analyst whose goal is to develop systems or procedures to support the pilot. Recent research which we have conducted on pilot decision making suggests that the pilot will seek to develop a single, concrete representation of the situation, typically a "worst case" scenario, but under some circumstances, a "best case" scenario. The analyst, on the other hand, will develop a system which mathematically aggregates the possibilities. An expected value will be computed for each option: i.e., a weighted average of possibilities, in which the probabilities assigned to each possible outcome serve as the weights. He thus provides an abstract level of representation (e.g., a display of "expected danger" contours) that corresponds to no actually realizable state of affairs.

A large literature now exists detailing biases and fallacies in unaided human reasoning about uncertainty. One theme in this work is that people often do not properly consider multiple possibilities, i.e., they suppress uncertainty. This literature is frequently cited to justify the introduction of computerized decision aids. Despite obligatory references to the notion of "supporting" the user rather than "replacing" him, the almost universal approach to aiding involves replacement of the human's preferred problem-solving method with the approach dictated by a normative model. A question that has seldom been asked is what (if any) *benefits* might be served by the human knowledge representations and decision-making strategies which are associated with biases.

Pilots do not want to become Bayesians. In our research, they strongly preferred single possibility displays (e.g., worst case) to probabilistically aggregated ones. Moreover, interviews with the pilots revealed a rather sophisticated and active process of problem solving underlying the selection (and rejection) of single possibility representations. Pilots adopted different assumptions depending on whether the uncertainty pertained to the existence of a threat or its classification, and whether uncertainty arose from incompleteness of data or conflicting evidence. When evidence is conflicting, pilots search for an explanation of the conflict and seek to resolve it by revising assumptions about the sources of data. It should not be particularly surprising if decision aids which are incompatible with the pilot's preferred way of thinking about the problem were simply rejected by their intended users. It would also not be surprising if errors resulted from an inappropriate interpretation of the aid's conclusions or method of reasoning. It is at this point that the analyst may take a strong stand: either train the users in the normatively correct (probabilistic) procedure, or automate. The pilot will not, of course, be happy with either of these alternatives: when decisions affect his survival or mission success, he insists on being in the loop and believes that he often has something to contribute to the solution.

A defense of the pilot may be at hand in recent theoretical research on knowledge representation. The human ability to generate and test hypotheses (e.g., explanation of unexpected events or actions to deal with unanticipated obstacles) may depend on mental models which represent the causal relationships among objects in the world in an essentially "analogical" fashion. The components of such models correspond one to one with represented objects in the world, and conclusions are "read off" the model itself without the benefit of previously existing general rules or knowledge. While analogical models may be the *sine qua non* of generating new knowledge, they are unable to represent indeterminacy or ambiguity effectively. The pilot's cognitive biases, therefore, may be inextricably intertwined with rather powerful information processing capabilities, which sometimes enable him to "think ahead of the airplane," solve novel problems, and handle unanticipated types of evidence. A decision aid which imposes an alien mode of thinking may fail to tap effectively user knowledge that could enhance overall system performance.

These considerations have led us to the concept of "personalized and prescriptive decision aiding." Traditional decision aids may throw out the baby (i.e., user knowledge) with the bath water (i.e., user biases). The aim of personalized and prescriptive decision aiding is to facilitate the user's preferred approach to the problem, and to

explore the possibility of a more precise, "surgical" removal of biases—by addressing them in the context of the decision maker's preferred approach. The "biases," of course, often look different when viewed in this context. For example, the pilot who adopts a worst case strategy is not really suppressing uncertainty; he knows perfectly well that other outcomes are possible. Rather, he is adopting assumptions which enable him to focus on a concrete, causally modeled state of affairs as opposed to an abstract, non-realizable average or expected value. He may subsequently wish to undo these particular assumptions and explore another set, which implies another concrete, causally modeled state of affairs. "Reducing biases" might, therefore, be better regarded as a matter of helping users keep track of their assumptions and alerting them at signs of trouble or when other assumptions would have substantially different implications.

In the case of our pilot, for example, this takes the form of (a) providing him single possibility (e.g., worst case) displays automatically; (b) giving him the option of requesting other single possibility (e.g., best case) displays; (c) prompting him when a promising action alternative has been overlooked (i.e., slightly more risky on worst case assumptions, but with significantly greater potential in other situations); (d) providing a simple graphical display of sources of data and the hypotheses they support; (e) alerting when sources of data are in conflict; and (f) providing both automated and interactive methods for exploring alternative assumptions regarding those sources. This type of decision aid works with the decision maker: rather than requiring him to think in abstract "normative" terms (i.e., to compare options in terms of their expected values), displays are provided which mesh naturally with his original approach (look only at single, concrete possibilities), but which support and expand that approach by facilitating the exploration of *multiple* possibilities where they matter.

Personalized and prescriptive decision-aiding is in no sense a matter of "slumming," i.e., adopting a less than optimal decision-making approach in deference to the user's preferences or cognitive limitations. Overall system performance with aids of this type should come very nearly as close to satisfying the normative constraints of "coherence" as do straightforward Bayesian aids. The Bayesian analyst should be pacified by observing that payoffs are often a "flat maximum" with respect to many features of the "normatively correct" solution, and that cases where significant suboptimality occurs can often be identified rather precisely (and appropriate prompts provided for the user). In a sense, though, this is too superficial a view of the situation. In fact, the Bayesian is in the same boat (rather, airplane) as the pilot. Probabilistic arguments, even though their conclusions are in the form of probabilities or degrees of belief, *also* depend on assumptions, and thus suppress uncertainty. These assumptions, which may be explicit or implicit, pertain to modeling (e.g., normality, independence, linearity) and also to substance (e.g., the credibility of a source, proper functioning of a data collection system, continued accuracy of dated observation, absence of a conspiracy to deceive among apparently unrelated sources). Unless such assumptions are made, human reasoning—whether the pilot's or the Bayesian's—is condemned to perpetual inconclusiveness. The essence of normative behavior is to keep track of the most problematic assumptions, to test the sensitivity of conclusions to changes in such assumptions, and to remain alert to signs of trouble; dispensing with assumptions altogether is not possible. Inference mechanisms and decision aids which embody this hypothesis-testing attitude are all too rare.

Our hypothesis is that displays which represent information in accordance with a user's own internal representations should be more readily utilized, should be understood more quickly and accurately, and should provide a more effective context for eliciting on-the-spot user knowledge.

Research Issues.

The success of the decision aid design methodology which we have recommended obviously depends on the resolution of a huge number of research issues. We see great value in basic research which would illuminate: the knowledge structures and decision-making strategies utilized in various domains and their linkage to biases in reasoning; the effectiveness of different methods for "reducing biases" which are compatible in spirit with user-preferred problem-solving strategies; methods for eliciting knowledge from potential users and mapping that knowledge onto appropriate displays; and finally, the exploration of richer normative models which more adequately capture the sophisticated processes of adopting and revising assumptions in which real decision makers engage.

Bibliography

Gentner, D., and Stevens, A.L. (Eds.). (1983). Mental models. Hillsdale, NJ: Erlbaum.

Johnson-Laird, P.N. (1983). Mental models. Cambridge, MA: Harvard University Press.

Kahneman, D., Slovic, P., and Tversky, A. (Eds.). (1982). Judgment under uncertainty: Heuristics and biases. NY: Cambridge University Press.

Report of the Research Briefing Panel on Decision Making and Problem Solving

Herbert Simon, Chair

(Reprinted with permission from "Research Briefings 1986,"
© 1986 by the National Academy of Sciences.)

INTRODUCTION

The work of managers, of scientists, of engineers, of lawyers—the work that steers the course of society and its economic and governmental organizations—is largely work of making decisions and solving problems. It is work of choosing issues that require attention, setting goals, finding or designing suitable courses of action, and evaluating and choosing among alternative actions. The first three of these activities—fixing agendas, setting goals, and designing actions—are usually called problem solving; the last, evaluating and choosing, is usually called decision making. Nothing is more important for the well-being of society than that this work be performed effectively, that we address successfully the many problems requiring attention at the national level (the budget and trade deficits, AIDS, national security, the mitigation of earthquake damage), at the level of business organizations (product improvement, efficiency of production, choice of investments), and at the level of our individual lives (choosing a career or a school, buying a house).

The abilities and skills that determine the quality of our decisions and problem solutions are stored not only in more than 200 million human heads, but also in tools and machines, and especially today in those machines we call computers. This fund of brains and its attendant machines form the basis of our American ingenuity, an ingenuity that has permitted U.S. society to reach remarkable levels of economic productivity.

There are no more promising or important targets for basic scientific research than understanding how human minds, with and without the help of computers, solve problems and make decisions effectively, and improving our problem-solving and decision-making capabilities. In psychology, economics, mathematical statistics, operations research, political science, artificial intelligence, and cognitive science, major research gains have been made during the past half century in understanding problem solving and decision making. The progress already achieved holds forth the promise of exciting new advances that will contribute substantially to our nation's capacity for dealing intelligently with the range of issues, large and small, that confront us.

Much of our existing knowledge about decision making and problem solving, derived from this research, has already been put to use in a wide variety of applications, including procedures used to assess drug safety, inventory control methods for industry, the new expert systems that embody artificial intelligence techniques, procedures for modeling energy and environmental systems, and analyses of the stabilizing or destabilizing effects of alternative defense strategies. (Application of the new inventory control techniques, for example, has enabled American corporations to reduce their inventories by hundreds of millions of dollars since World War II without increasing the incidence of stockouts.) Some of the knowledge gained through the research describes the ways in which people actually go about making decisions and solving problems; some of it prescribes better methods, offering advice for the improvement of the process.

Central to the body of prescriptive knowledge about decision making has been the theory of subjective expected utility (SEU), a sophisticated mathematical model of choice that lies at the foundation of most contemporary economics, theoretical statistics, and operations research. SEU theory defines the conditions of perfect utility-maximizing rationality in a world of certainty or in a world in which the probability distributions of all relevant variables can be provided by the decision makers. (In spirit, it might be compared with a theory of ideal gases or of frictionless bodies sliding down inclined planes in a vacuum.) SEU theory deals only with decision making; it has nothing to say about how to frame problems, set goals, or develop new alternatives.

Prescriptive theories of choice such as SEU are complemented by empirical research that shows how people actually make decisions (purchasing insurance, voting for political candidates, or investing in securities), and research on the processes people use to solve problems (designing switchgear or finding chemical reaction pathways). This research demonstrates that people solve problems by selective, heuristic search through large problem spaces and large data bases, using means-ends analysis as a principal technique for guiding the search. The expert systems that are now being produced by research on artificial intelligence and applied to such tasks as interpreting oil-well drilling logs or making medical diagnoses are outgrowths of these research findings on human problem solving.

What chiefly distinguishes the empirical research on decision making and problem solving from the prescriptive approaches derived from SEU theory is the attention that the former gives to the limits on human rationality. These limits are imposed by the complexity of the world in which we live, the incompleteness and inadequacy of human knowledge, the inconsistencies of individual preference and belief, the conflicts of value among people and groups of people, and the inadequacy of the computations we can carry out, even with the aid of the most powerful computers. The real world of human decisions is not a world of ideal gases, frictionless planes, or vacuums. To bring it within the scope of human thinking powers, we must simplify our problem formulations drastically, even leaving out much or most of what is potentially relevant.

The descriptive theory of problem solving and decision making is centrally concerned with how people cut problems down to size: how they apply approximate, heuristic techniques to handle complexity that cannot be handled exactly. Out of this descriptive theory is emerging an augmented and amended prescriptive theory, one that takes account of the gaps and elements of unrealism in SEU theory by encompassing problem solving as well as choice and demanding only the kinds of knowledge, consistency, and computational power that are attainable in the real world.

The growing realization that coping with complexity is central to human decision making strongly influences the directions of research in this domain. Operations research and artificial intelligence are forging powerful new computational tools; at the same time, a new body of mathematical theory is evolving around the topic of computational complexity. Economics, which has traditionally derived both its descriptive and prescriptive approaches from SEU theory, is now paying a great deal of attention to uncertainty and incomplete information; to so-called "agency theory," which takes account of the institutional framework within which decisions are made; and to game theory, which seeks to deal with interindividual and intergroup processes in which there is partial conflict of interest. Economists and political scientists are also increasingly buttressing the empirical foundations of their field by studying individual choice behavior directly and by studying behavior in experimentally constructed markets and simulated political structures.

The following pages contain a fuller outline of current knowledge about decision making and problem solving and a brief review of current research directions in these fields as well as some of the principal research opportunities.

DECISION MAKING

SEU Theory

The development of SEU theory was a major intellectual achievement of the first half of this century. It gave for the first time a formally axiomatized statement of what it would mean for an agent to behave in a consistent, rational manner. It assumed that a decision maker possessed a utility function (an ordering by preference among all the possible outcomes of choice), that all the alternatives among which choice could be made were known, and that the consequences of choosing each alternative could be ascertained (or, in the version of the theory that treats choice under uncertainty, it assumed that a subjective or objective probability distribution of conse-

quences was associated with each alternative). By admitting subjectively assigned probabilities, SEU theory opened the way to fusing subjective opinions with objective data, an approach that can also be used in man-machine decision-making systems. In the probabilistic version of the theory, Bayes's rule prescribes how people should take account of new information and how they should respond to incomplete information.

The assumptions of SEU theory are very strong, permitting correspondingly strong inferences to be made from them. Although the assumptions cannot be satisfied even remotely for most complex situations in the real world, they may be satisfied approximately in some microcosms—problem situations that can be isolated from the world's complexity and dealt with independently. For example, the manager of a commercial cattle-feeding operation might isolate the problem of finding the least expensive mix of feeds available in the market that would meet all the nutritional requirements of his cattle. The computational tool of linear programming, which is a powerful method for maximizing goal achievement or minimizing costs while satisfying all kinds of side conditions (in this case, the nutritional requirements), can provide the manager with an optimal feed mix—optimal within the limits of approximation of his model to real-world conditions. Linear programming and related operations research techniques are now used widely to make decisions whenever a situation that reasonably fits their assumptions can be carved out of its complex surround. These techniques have been especially valuable aids to middle management in dealing with relatively well-structured decision problems.

Most of the tools of modern operations research—not only linear programming, but also integer programming, queuing theory, decision trees, and other widely used techniques—use the assumptions of SEU theory. They assume that what is desired is to maximize the achievement of some goal, under specified constraints and assuming that all alternatives and consequences (or their probability distributions) are known. These tools have proven their usefulness in a wide variety of applications.

The Limits of Rationality

Operations research tools have also underscored dramatically the limits of SEU theory in dealing with complexity. For example, present and prospective computers are not even powerful enough to provide exact solutions for the problems of optimal scheduling and routing of jobs through a typical factory that manufactures a variety of products using many different tools and machines. And the mere thought of using these computational techniques to determine an optimal national policy for energy production or an optimal economic policy reveals their limits.

Computational complexity is not the only factor that limits the literal application of SEU theory. The theory also makes enormous demands on information. For the utility function, the range of available alternatives and the consequences following from each alternative must all be known. Increasingly, research is being directed at decision making that takes realistic account of the compromises and approximations that must be made in order to fit real-world problems to the informational and computational limits of people and computers, as well as to the inconsistencies in their values and perceptions. The study of actual decision processes (for example, the strategies used by corporations to make their investments) reveals massive and unavoidable departures from the framework of SEU theory. The sections that follow describe some of the things that have been learned about choice under various conditions of incomplete information, limited computing power, inconsistency, and institutional constraints on alternatives. Game theory, agency theory, choice under uncertainty, and the theory of markets are a few of the directions of this research, with the aims both of constructing prescriptive theories of broader application and of providing more realistic descriptions and explanations of actual decision making within U.S. economic and political institutions.

Limited Rationality in Economic Theory

Although the limits of human rationality were stressed by some researchers in the 1950's, only recently has there been extensive activity in the field of economics aimed at developing theories that assume less than fully rational choice on the part of business firm managers and other economic agents. The newer theoretical research

undertakes to answer such questions as the following:

- o Are market equilibria altered by the departures of actual choice behavior from the behavior of fully rational agents predicted by SEU theory?
- o Under what circumstances do the processes of competition "police" markets in such a way as to cancel out the effects of the departures from full rationality?
- o In what ways are the choices made by boundedly rational agents differ from those made by fully rational agents?

Theories of the firm that assume managers are aiming at "satisfactory" profits or that their concern is to maintain the firm's share of market in the industry make quite different predictions about economic equilibrium than those derived from the assumption of profit maximization. Moreover, the classical theory of the firm cannot explain why economic activity is sometimes organized around large business firms and sometimes around contractual networks of individuals or smaller organizations. New theories that take account of differential access of economic agents to information, combined with differences in self-interest, are able to account for these important phenomena, as well as provide explanations for the many forms of contracts that are used in business. Incompleteness and asymmetry of information have been shown to be essential for explaining how individuals and business firms decide when to face uncertainty by insuring, when by hedging, and when by assuming the risk.

Most current work in this domain still assumes that economic agents seek to maximize utility, but within limits posed by the incompleteness and uncertainty of the information available to them. An important potential area of research is to discover how choices will be changed if there are other departures from the axioms of rational choice—for example, substituting goals of reaching specified aspiration levels (satisficing) for goals of maximizing.

Applying the new assumptions about choice to economics leads to new empirically supported theories about decision making over time. The classical theory of perfect rationality leaves no room for regrets, second thoughts, or "weakness of will." It cannot explain why many individuals enroll in Christmas savings plans, which earn interest well below the market rate. More generally, it does not lead to correct conclusions about the important social issues of saving and conservation. The effect of pensions and social security on personal saving has been a controversial issue in economics. The standard economic model predicts that an increase in required pension saving will reduce other saving dollar for dollar; behavioral theories, on the other hand, predict a much smaller offset. The empirical evidence indicates that the offset is indeed very small. Another empirical finding is that the method of payment of wages and salaries affects the saving rate. For example, annual bonuses produce a higher saving rate than the same amount of income paid in monthly salaries. This finding implies that saving rates can be influenced by the way compensation is framed.

If individuals fail to discount properly for the passage of time, their decisions will not be optimal. For example, air conditioners vary greatly in their energy efficiency; the more efficient models cost more initially but save money over the long run through lower energy consumption. It has been found that consumers, on average, choose air conditioners that imply a discount rate of 25 percent or more per year, much higher than the rates of interest that prevailed at the time of the study.

As recently as five years ago, the evidence was thought to be unassailable that markets like the New York Stock Exchange work efficiently—that prices reflect all available information at any given moment in time, so that stock price movements resemble a random walk and contain no systematic information that could be exploited for profit. Recently, however, substantial departures from the behavior predicted by the efficient market hypothesis have been detected. For example, small firms appear to earn inexplicably high returns on the market prices of their stock, while firms that have very low price-earnings ratios and firms that have lost much of their market value in the recent past also earn abnormally high returns. All of these results are consistent with the empirical finding that decision makers often overreact to new information, in violation of Bayes's rules. In the same way, it has been found that stock prices are excessively volatile—that they fluctuate up and down more rapidly and violently than they would if the market were efficient.

There has also been a long-standing puzzle as to why firms pay dividends. Considering that dividends are taxed at a higher rate than capital gains, taxpaying investors should prefer, under the assumptions of perfect rationality, that their firms reinvest earnings or repurchase shares instead of paying dividends. (The investors could simply sell some of their appreciated shares to obtain the income they require.) The solution to this puzzle also requires models of investors that take account of limits on rationality.

The Theory of Games

In economic, political, and other social situations in which there is actual or potential conflict of interest, especially if it is combined with incomplete information, SEU theory faces special difficulties. In markets in which there are many competitors (e.g., the wheat market), each buyer or seller can accept the market price as a "given" that will not be affected materially by the actions of any single individual. Under these conditions, SEU theory makes unambiguous predictions of behavior. However, when a market has only a few suppliers—say, for example, two—matters are quite different. In this case, what it is rational to do depends on what one's competitor is going to do, and vice versa. Each supplier may try to outwit the other. What then is the rational decision?

The most ambitious attempt to answer questions of this kind was the theory of games, developed by von Neumann and Morgenstern and published in its full form in 1944. But the answers provided by the theory of games are sometimes very puzzling and ambiguous. In many situations, no single course of action dominates all the others; instead, a whole set of possible solutions are all equally consistent with the postulates of rationality.

One game that has been studied extensively, both theoretically and empirically, is the Prisoner's Dilemma. In this game between two players, each has a choice between two actions, one trustful of the other player, the other mistrustful or exploitative. If both players choose the trustful alternative, both receive small rewards. If both choose the exploitative alternative, both are punished. If one chooses the trustful alternative and the other the exploitative alternative, the former is punished much more severely than in the previous case, while the latter receives a substantial reward. If the other player's choice is fixed but unknown, it is advantageous for a player to choose the exploitative alternative, for this will give him the best outcome in either case. But if both adopt this reasoning, they will both be punished, whereas they could both receive rewards if they agreed upon the trustful choice (and did not welch on the agreement).

The terms of the game have an unsettling resemblance to certain situations in the relations between nations or between a company and the employees' union. The resemblance becomes stronger if one imagines the game as being played repeatedly. Analyses of "rational" behavior under assumptions of intended utility maximization support the conclusion that the players will (ought to?) always make the mistrustful choice. Nevertheless, in laboratory experiments with the game, it is often found that players (even those who are expert in game theory) adopt a "tit-for-tat" strategy. That is, each plays the trustful, cooperative strategy as long as his or her partner does the same. If the partner exploits the player on a particular trial, the player then plays the exploitative strategy on the next trial and continues to do so until the partner switches back to the trustful strategy. Under these conditions, the game frequently stabilizes with the players pursuing the mutually trustful strategy and receiving the rewards.

With these empirical findings in hand, theorists have recently sought and found some of the conditions for attaining this kind of benign stability. It occurs, for example, if the players set aspirations for a satisfactory reward rather than seeking the maximum reward. This result is consistent with the finding that in many situations, as in the Prisoner's Dilemma game, people appear to satisfice rather than attempting to optimize.

The Prisoner's Dilemma game illustrates an important point that is beginning to be appreciated by those who do research on decision making. There are so many ways in which actual human behavior can depart from the SEU assumptions that theorists seeking to account for behavior are confronted with an embarrassment of riches. To choose among the many alternative models that could account for the anomalies of choice, extensive empirical research is called for—to see how people do make their choices, what beliefs guide them, what information they have available, and what part of that information they take into account and what part they ignore. In a world of

limited rationality, economics and the other decision sciences must closely examine the actual limits on rationality in order to make accurate predictions and to provide sound advice on public policy.

Empirical Studies of Choice Under Uncertainty

During the past 10 years, empirical studies of human choices in which uncertainty, inconsistency, and incomplete information are present have produced a rich collection of findings which only now are beginning to be organized under broad generalizations. Here are a few examples. When people are given information about the probabilities of certain events (e.g., how many lawyers and how many engineers are in a population that is being sampled), and then are given some additional information as to which of the events has occurred (which person has been sampled from the population), they tend to ignore the prior probabilities in favor of incomplete or even quite irrelevant information about the individual event. Thus, if they are told that 70 percent of the population are lawyers, and if they are then given a noncommittal description of a person (one that could equally well fit a lawyer or an engineer), half the time they will predict that the person is a lawyer and half the time that he is an engineer—even though the laws of probability dictate that the best forecast is always to predict that the person is a lawyer.

People commonly misjudge probabilities in many other ways. Asked to estimate the probability that 60 percent or more of the babies born in a hospital during a given week are male, they ignore information about the total number of births, although it is evident that the probability of a departure of this magnitude from the expected value of 50 percent is smaller if the total number of births is larger (the standard error of a percentage varies inversely with the square root of the population size).

There are situations in which people assess the frequency of a class by the ease with which instances can be brought to mind. In one experiment, subjects heard a list of names of persons of both sexes and were later asked to judge whether there were more names of men or women on the list. In lists presented to some subjects, the men were more famous than the women; in other lists, the women were more famous than the men. For all lists, subjects judged that the sex that had the more famous personalities was the more numerous.

The way in which an uncertain possibility is presented may have a substantial effect on how people respond to it. When asked whether they would choose surgery in a hypothetical medical emergency, many more people said that they would when the chance of survival was given as 80 percent than when the chance of death was given as 20 percent.

On the basis of these studies, some of the general heuristics, or rules of thumb, that people use in making judgments have been compiled—heuristics that produce biases toward classifying situations according to their representativeness, or toward judging frequencies according to the availability of examples in memory, or toward interpretations warped by the way in which a problem has been framed. These findings have important implications for public policy. A recent example is the lobbying effort of the credit card industry to have differentials between cash and credit prices labeled “cash discounts” rather than “credit surcharges.” The research findings raise questions about how to phrase cigarette warning labels or frame truth-in-lending laws and informed consent laws.

Methods of Empirical Research

Finding the underlying bases of human choice behavior is difficult. People cannot always, or perhaps even usually, provide veridical accounts of how they make up their minds, especially when there is uncertainty. In many cases, they can predict how they will behave (pre-election polls of voting intentions have been reasonably accurate when carefully taken), but the reasons people give for their choices can often be shown to be rationalizations and not truly related to their real motives.

Students of choice behavior have steadily improved their research methods. They question respondents about specific situations, rather than asking for generalizations. They are sensitive to the dependence of answers on the exact forms of the questions. They are aware that behavior in an experimental situation may be different from

behavior in real life, and they attempt to provide experimental settings and motivations that are as realistic as possible. Using thinking-aloud protocols and other approaches, they try to track the choice behavior step by step, instead of relying just on information about outcomes or querying respondents retrospectively about their choice processes.

Perhaps the most common method of empirical research in this field is still to ask people to respond to a series of questions. But data obtained by this method are being supplemented by data obtained from carefully designed laboratory experiments and from observations of actual choice behavior (for example, the behavior of customers in supermarkets). In an experimental study of choice, subjects may trade in an actual market with real (if modest) monetary rewards and penalties. Research experience has also demonstrated the feasibility of making direct observations, over substantial periods of time, of the decision-making processes in business and governmental organizations—for example, observations of the procedures that corporations use in making new investments in plant and equipment. Confidence in the empirical findings that have been accumulating over the past several decades is enhanced by the general consistency that is observed among the data obtained from quite different settings using different research methods.

There still remains the enormous and challenging task of putting together these findings into an empirically founded theory of decision making. With the growing availability of data, the theory-building enterprise is receiving much better guidance from the facts than it did in the past. As a result, we can expect it to become correspondingly more effective in arriving at realistic models of behavior.

PROBLEM SOLVING

The theory of choice has its roots mainly in economics, statistics, and operations research and only recently has received much attention from psychologists; the theory of problem solving has a very different history. Problem solving was initially studied principally by psychologists, and more recently by researchers in artificial intelligence. It has received rather scant attention from economists.

Contemporary Problem-Solving Theory

Human problem solving is usually studied in laboratory settings, using problems that can be solved in relatively short periods of time (seldom more than an hour), and often seeking a maximum density of data about the solution process by asking subjects to think aloud while they work. The thinking-aloud technique, at first viewed with suspicion by behaviorists as subjective and "introspective," has received such careful methodological attention in recent years that it can now be used dependably to obtain data about subjects' behaviors in a wide range of settings.

The laboratory study of problem solving has been supplemented by field studies of professionals solving real-world problems—for example, physicians making diagnoses and chess grandmasters analyzing game positions, and, as noted earlier, even business corporations making investment decisions. Currently, historical records, including laboratory notebooks of scientists, are also being used to study problem-solving processes in scientific discovery. Although such records are far less "dense" than laboratory protocols, they sometimes permit the course of discovery to be traced in considerable detail. Laboratory notebooks of scientists as distinguished as Charles Darwin, Michael Faraday, Antoine-Laurent Lavoisier, and Hans Krebs have been used successfully in such research.

From empirical studies, a description can now be given of the problem-solving process that holds for a rather wide range of activities. First, problem solving generally proceeds by selective search through large sets of possibilities, using rules of thumb (heuristics) to guide the search. Because the possibilities in realistic problem situations are generally multitudinous, trial-and-error search would simply not work; the search must be highly selective. Chess grandmasters seldom examine more than a hundred of the vast number of possible scenarios that confront them, and similar small numbers of searches are observed in other kinds of problem-solving search.

One of the procedures often used to guide search is "hill-climbing," using some measure of approach to the goal to determine where it is most profitable to look next. Another, and more powerful, common procedure is means-ends analysis. In means-ends analysis, the problem solver compares the present situation with the goal, detects a difference between them, and then searches memory for actions that are likely to reduce the difference. Thus, if the difference is a 50-mile distance from the goal, the problem solver will retrieve from memory knowledge about autos, carts, bicycles, and other means of transport; walking and flying will probably be discarded as inappropriate for that distance.

The third thing that has been learned about problem solving—especially when the solver is an expert—is that it relies on large amounts of information that are stored in memory and that are retrievable whenever the solver recognizes cues signaling its relevance. Thus, the expert knowledge of a diagnostician is evoked by the symptoms presented by the patient; this knowledge leads to the recollection of what additional information is needed to discriminate among alternative diseases and, finally, to the diagnosis.

In a few cases, it has been possible to estimate how many patterns an expert must be able to recognize in order to gain access to the relevant knowledge stored in memory. A chess master must be able to recognize about 50,000 different configurations of chess pieces that occur frequently in the course of chess games. A medical diagnostician must be able to recognize tens of thousands of configurations of symptoms; a botanist or zoologist specializing in taxonomy, tens or hundreds of thousands of features of specimens that define their species. For comparison, college graduates typically have vocabularies in their native languages of 50,000 to 200,000 words. (However, these numbers are very small in comparison with the real-world situations the expert faces: there are perhaps 10^{120} branches in the game tree of chess, a game played with only six kinds of pieces on an 8X8 board.)

One of the accomplishments of the contemporary theory of problem solving has been to provide an explanation for the phenomena of intuition and judgment frequently seen in experts' behavior. The store of expert knowledge, "indexed" by the recognition cues that make it accessible and combined with some basic inferential capabilities (perhaps in the form of means-ends analysis), accounts for the ability of experts to find satisfactory solutions for difficult problems, and sometimes to find them almost instantaneously. The expert's "intuition" and "judgment" derive from this capability for rapid recognition linked to a large store of knowledge. When immediate intuition fails to yield a problem solution or when a prospective solution needs to be evaluated, the expert falls back on the slower processes of analysis and inference.

Expert Systems in Artificial Intelligence

Over the past 30 years, there has been close teamwork between research in psychology and research in computer science aimed at developing intelligent programs. Artificial intelligence (AI) research has both borrowed from and contributed to research on human problem solving. Today, artificial intelligence is beginning to produce systems, applied to a variety of tasks, that can solve difficult problems at the level of professionally trained humans. These AI programs are usually called expert systems. A description of a typical expert system would resemble closely the description given above of typical human problem solving; the differences between the two would be differences in degree, not in kind. An AI expert system, relying on the speed of computers and their ability to retain large bodies of transient information in memory, will generally use "brute force"—sheer computational speed and power—more freely than a human expert can. A human expert, in compensation, will generally have a richer set of heuristics to guide search and a larger vocabulary of recognizable patterns. To the observer, the computer's process will appear the more systematic and even compulsive, the human's the more intuitive. But these are quantitative, not qualitative, differences.

The number of tasks for which expert systems have been built is increasing rapidly. One is medical diagnosis (two examples are the CADUCEUS and MYCIN programs). Others are automatic design of electric motors, generators, and transformers (which predates by a decade the invention of the term "expert systems"), the configuration of computer systems from customer specifications, and the automatic generation of reaction paths for the synthesis of organic molecules. All of these (and others) are either being used currently in professional or industrial practice or at least have reached a level at which they can produce a professionally acceptable product.

Expert systems are generally constructed in close consultation with the people who are experts in the task domain. Using standard techniques of observation and interrogation, the heuristics that the human expert uses, implicitly and often unconsciously, to perform the task are gradually educed, made explicit, and incorporated in program structures. Although a great deal has been learned about how to do this, improving techniques for designing expert systems is an important current direction of research. It is especially important because expert systems, once built, cannot remain static but must be modifiable to incorporate new knowledge as it becomes available.

Dealing With Ill-Structured Problems

In the 1950s, and 1960s, research on problem solving focused on clearly structured puzzle-like problems that were easily brought into the psychological laboratory and that were within the range of computer programming sophistication at that time. Computer programs were written to discover proofs for theorems in Euclidean geometry or to solve the puzzle of transporting missionaries and cannibals across a river. Choosing chess moves was perhaps the most complex task that received attention in the early years of cognitive science and AI.

As understanding grew of the methods needed to handle these relatively simple tasks, research aspirations rose. The next main target, in the 1960s and 1970s, was to find methods for solving problems that involved large bodies of semantic information. Medical diagnosis and interpreting mass spectrogram data are examples of the kinds of tasks that were investigated during this period and for which a good level of understanding was achieved. They are tasks that, for all of the knowledge they call upon, are still well structured, with clear-cut goals and constraints.

The current research target is to gain an understanding of problem-solving tasks when the goals themselves are complex and sometimes ill defined, and when the very nature of the problem is successively transformed in the course of exploration. To the extent that a problem has these characteristics, it is usually called ill structured. Because ambiguous goals and shifting problem formulations are typical characteristics of problems of design, the work of architects offers a good example of what is involved in solving ill-structured problems. An architect begins with some very general specifications of what is wanted by a client. The initial goals are modified and substantially elaborated as the architect proceeds with the task. Initial design ideas, recorded in drawings and diagrams, themselves suggest new criteria, new possibilities, and new requirements. Throughout the whole process of design, the emerging conception provides continual feedback that reminds the architect of additional considerations that need to be taken into account.

With the current state of the art, it is just beginning to be possible to construct programs that simulate this kind of flexible problem-solving process. What is called for is an expert system whose expertise includes substantial knowledge about design criteria as well as knowledge about the means for satisfying those criteria. Both kinds of knowledge are evoked in the course of the design activity by the usual recognition processes, and the evocation of design criteria and constraints continually modifies and remolds the problem that the design system is addressing. The large data bases that can now be constructed to aid in the management of architectural and construction projects provide a framework into which AI tools, fashioned along these lines, can be incorporated.

Most corporate strategy problems and governmental policy problems are at least as ill structured as problems of architectural or engineering design. The tools now being forged for aiding architectural design will provide a basis for building tools that can aid in formulating, assessing, and monitoring public energy or environmental policies, or in guiding corporate product and investment strategies.

Setting the Agenda and Representing a Problem

The very first steps in the problem-solving process are the least understood. What brings (and should bring) problems to the head of the agenda? And when a problem is identified, how can it be represented in a way that facilitates its solution?

The task of setting an agenda is of utmost importance because both individual human beings and human institutions have limited capacities for dealing with many tasks simultaneously. While some problems are receiving full attention, others are neglected. When new problems come thick and fast, "fire fighting" replaces planning and deliberation. The facts of limited attention span, both for individuals and for institutions like the Congress, are well known. However, relatively little has been accomplished toward analyzing or designing effective agenda-setting systems. A beginning could be made by the study of "alerting" organizations like the Office of Technology Assessment or military and foreign affairs intelligence agencies. Because the research and development function in industry is also in considerable part a task of monitoring current and prospective technological advances, it could also be studied profitably from this standpoint.

The way in which problems are represented has much to do with the quality of the solutions that are found. The task of designing highways or dams takes on an entirely new aspect if human responses to a changed environment are taken into account. (New transportation routes cause people to move their homes, and people show a considerable propensity to move into zones that are subject to flooding when partial protections are erected.) Very different social welfare policies are usually proposed in response to the problem of providing incentives for economic independence than are proposed in response to the problem of taking care of the needy. Early management information systems were designed on the assumption that information was the scarce resource; today, because designers recognize that the scarce resource is managerial attention, a new framework produces quite different designs.

The representation or "framing" of problems is even less well understood than agenda setting. Today's expert systems make use of problem representations that already exist. But major advances in human knowledge frequently derive from new ways of thinking about problems. A large part of the history of physics in nineteenth-century England can be written in terms of the shift from action-at-a-distance representations to the field representations that were developed by the applied mathematicians at Cambridge.

Today, developments in computer-aided design (CAD) present new opportunities to provide human designers with computer-generated representations of their problems. Effective use of these capabilities requires us to understand better how people extract information from diagrams and other displays and how displays can enhance human performance in design tasks. Research on representations is fundamental to the progress of CAD.

Computation as Problem Solving

Nothing has been said so far about the radical changes that have been brought about in problem solving over most of the domains of science and engineering by the standard uses of computers as computational devices. Although a few examples come to mind in which artificial intelligence has contributed to these developments, they have mainly been brought about by research in the individual sciences themselves, combined with work in numerical analysis.

Whatever their origins, the massive computational applications of computers are changing the conduct of science in numerous ways. There are new specialities emerging such as "computational physics" and "computational chemistry." Computation—that is to say, problem solving—becomes an object of explicit concern to scientists, side by side with the substance of the science itself. Out of this new awareness of the computational component of scientific inquiry is arising an increasing interaction among computational specialists concerned with cognition and AI. This interaction extends well beyond the traditional area of numerical analysis, or even the newer subject of computational complexity, into the heart of the theory of problem solving.

Physicists seeking to handle the great mass of bubble-chamber data produced by their instruments began, as early as the 1960s, to look to AI for pattern recognition methods as a basis for automating the analysis of their data. The construction of expert systems to interpret mass spectrogram data and of other systems to design synthesis paths for chemical reactions are other examples of problem solving in science, as are programs to aid in matching sequences of nucleic acids in DNA and RNA and amino acid sequences in proteins.

Theories of human problem solving and learning are also beginning to attract new attention within the scientific community as a basis for improving science teaching. Each advance in the understanding of problem solving and learning processes provides new insights about the ways in which a learner must store and index new knowledge and procedures if they are to be useful for solving problems. Research on these topics is also generating new ideas about how effective learning takes place—for example, how students can learn by examining and analyzing worked-out examples.

EXTENSIONS OF THEORY

Opportunities for advancing our understanding of decision making and problem solving are not limited to the topics dealt with above, and in this section, just a few indications of additional promising directions for research are presented.

Decision Making Over Time

The time dimension is especially troublesome in decision making. Economics has long used the notion of time discounting and interest rates to compare present with future consequences of decisions, but as noted above, research on actual decision making shows that people frequently are inconsistent in their choices between present and future. Although time discounting is a powerful idea, it requires fixing appropriate discount rates for individual, and especially social, decisions. Additional problems arise because human tastes and priorities change over time. Classical SEU theory assumes a fixed, consistent utility function, which does not easily accommodate changes in taste. At the other extreme, theories postulating a limited attention span do not have ready ways of ensuring consistency of choice over time.

Aggregation

In applying our knowledge of decision making and problem solving to society-wide, or even organization-wide, phenomena, the problem of aggregation must be solved; that is, ways must be found to extrapolate from theories of individual decision processes to the net effects on the whole economy, polity, and society. Because of the wide variety of ways in which any given decision task can be approached, it is unrealistic to postulate a "representative firm" or an "economic man," and to simply lump together the behaviors of large numbers of supposedly identical individuals. Solving the aggregation problem becomes more important as more of the empirical research effort is directed toward studying behavior at a detailed, microscopic level.

Organizations

Related to aggregation is the question of how decision making and problem solving change when attention turns from the behavior of isolated individuals to the behavior of these same individuals operating as members of organizations or other groups. When people assume organizational positions, they adapt their goals and values to their responsibilities. Moreover, their decisions are influenced substantially by the patterns of information flow and other communications among the various organization units.

Organizations sometimes display sophisticated capabilities far beyond the understanding of single individuals. They sometimes make enormous blunders or find themselves incapable of acting. Organizational performance is highly sensitive to the quality of the routines or "performance programs" that govern behavior and to the adaptability of these routines in the face of a changing environment. In particular, the "peripheral vision" of a complex organization is limited, so that responses to novelty in the environment may be made in inappropriate and quasi-automatic ways that cause major failure.

Theory development, formal modeling, laboratory experiments, and analysis of historical cases are all going forward in this important area of inquiry. Although the decision-making processes of organizations have been studied in the field on a limited scale, a great many more such intensive studies will be needed before the full range of techniques used by organizations to make their decisions is understood, and before the strengths and weaknesses of these techniques are grasped.

Learning

Until quite recently, most research in cognitive science and artificial intelligence had been aimed at understanding how intelligent systems perform their work. Only in the past five years has attention begun to turn to the question of how systems become intelligent—how they learn. A number of promising hypotheses about learning mechanisms are currently being explored. One is the so-called connexionist hypothesis, which postulates networks that learn by changing the strengths of their interconnections in response to feedback. Another learning mechanism that is being investigated is the adaptive production system, a computer program that learns by generating new instructions that are simply annexed to the existing program. Some success has been achieved in constructing adaptive production systems that can learn to solve equations in algebra and to do other tasks at comparable levels of difficulty.

Learning is of particular importance for successful adaptation to an environment that is changing rapidly. Because that is exactly the environment of the 1980s, the trend toward broadening research on decision making to include learning and adaptation is welcome.

This section has by no means exhausted the areas in which exciting and important research can be launched to deepen understanding of decision making and problem solving. But perhaps the examples that have been provided are sufficient to convey the promise and significance of this field of inquiry today.

CURRENT RESEARCH PROGRAMS

Most of the current research on decision making and problem solving is carried on in universities, frequently with the support of government funding agencies and private foundations. Some research is done by consulting firms in connection with their development and application of the tools of operations research, artificial intelligence, and systems modeling. In some cases, government agencies and corporations have supported the development of planning models to aid them in their policy planning—for example, corporate strategic planning for investments and markets and government planning of environmental and energy policies. There is an increasing number of cases in which research scientists are devoting substantial attention to improving the problem-solving and decision-making tools in their disciplines, as we noted in the examples of automation of the processing of bubble-chamber tracks and of the interpretation of mass spectrogram data.

To use a generous estimate, support for basic research in the areas described in this document is probably at the level of tens of millions of dollars per year, and almost certainly, it is not as much as \$100 million. The principal costs are for research personnel and computing equipment, the former being considerably larger.

Because of the interdisciplinary character of the research domain, federal research support comes from a number of different agencies, and it is not easy to assess the total picture. Within the National Science Foundation (NSF), the grants of the decision and management sciences, political science and the economics programs in the Social Sciences Division are to a considerable extent devoted to projects in this domain. Smaller amounts of support come from the memory and cognitive processes program in the Division of Behavioral and Neural Sciences, and perhaps from other programs. The "software" component of the new NSF Directorate of Computer Science and Engineering contains programs that have also provided important support to the study of decision making and problem solving.

The Office of Naval Research has, over the years, supported a wide range of studies of decision making, including important early support for operations research. The main source of funding for research in AI has been the Defense Advanced Research Projects Agency (DARPA) in the Department of Defense; important support for research on applications of AI to medicine has been provided by the National Institutes of Health.

Relevant economics research is also funded by other federal agencies; including the Treasury Department, the Bureau of Labor Statistics, and the Federal Reserve Board. In recent years, basic studies of decision making have received only relatively minor support from these sources, but because of the relevance of the research to their missions, they could become major sponsors.

Although a number of projects have been and are funded by private foundations, there appears to be at present no foundation for which decision making and problem solving are a major focus of interest.

In sum, the pattern of support for research in this field shows a healthy diversity but no agency with a clear lead responsibility, unless it be the rather modestly funded program in decision and management sciences at NSF. Perhaps the largest scale of support has been provided by DARPA, where decision making and problem solving are only components within the larger area of artificial intelligence and certainly not highly visible research targets.

The character of the funding requirements in this domain is much the same as in other fields of research. A rather intensive use of computational facilities is typical of most, but not all, of the research. And because the field is gaining new recognition and growing rapidly, there are special needs for the support of graduate students and postdoctoral training. In the computing-intensive part of the domain, desirable research funding per principal investigator might average \$250,000 per year; in empirical research involving field studies and large-scale experiments, a similar amount; and in other areas of theory and laboratory experimentation, somewhat less.

RESEARCH OPPORTUNITIES: SUMMARY

The study of decision making and problem solving has attracted much attention through most of this century. By the end of World War II, a powerful prescriptive theory of rationality, the theory of subjective expected utility (SEU), had taken form; it was followed by the theory of games. The past 40 years have seen widespread applications of these theories in economics, operations research, and statistics, and, through these disciplines, to decision making in business and government.

The main limitations of SEU theory and the developments based on it are its relative neglect of the limits of human (and computer) problem-solving capabilities in the face of real-world complexity. Recognition of these limitations has produced an increasing volume of empirical research aimed at discovering how humans cope with complexity and reconcile it with their bounded computational powers. Recognition that human rationality is limited occasions no surprise. What is surprising are some of the forms these limits take and the kinds of departures from the behavior predicted by the SEU model that have been observed. Extending empirical knowledge of actual human cognitive processes and of techniques for dealing with complexity continues to be a research goal of very high priority. Such empirical knowledge is needed both to build valid theories of how the U.S. society and economy operate and to build prescriptive tools for decision making that are compatible with existing computational capabilities.

The complementary fields of cognitive psychology and artificial intelligence have produced in the past 30 years a fairly well-developed theory of problem solving that lends itself well to computer simulation, both for purposes of testing its empirical validity and for augmenting human problem-solving capacities by the construction of expert systems. Problem-solving research today is being extended into the domain of ill-structured problems and applied to the task of formulating problem representations. The processes for setting the problem agenda, which are still very little explored, deserve more research attention.

The growing importance of computational techniques in all of the sciences has attracted new attention to numerical analysis and to the topic of computational complexity. The need to use heuristic as well as rigorous methods for analyzing very complex domains is beginning to bring about a wide interest, in various sciences, in the possible application of problem-solving theories to computation.

Opportunities abound for productive research in decision making and problem solving. A few of the directions of research that look especially promising and significant follow:

- o A substantially enlarged program of empirical studies, involving direct observation of behavior at the level of the individual and the organization, and including both laboratory and field experiments, will be essential in sifting the wheat from the chaff in the large body of theory that now exists and in giving direction to the development of new theory.
- o Expanded research on expert systems will require extensive empirical study of expert behavior and will provide a setting for basic research on how ill-structured problems are, and can be solved.
- o Decision making in organizational settings, which is much less well understood than individual decision making and problem solving, can be studied with great profit using already established methods of inquiry, especially through intensive long-range studies within individual organizations.
- o The resolution of conflicts of values (individual and group) and of inconsistencies in belief will continue to be highly productive directions of inquiry, addressed to issues of great importance to society.
- o Setting agendas and framing problems are two related but poorly understood processes that require special research attention and that now seem open to attack.

These five areas are examples of especially promising research opportunities drawn from the much larger set that are described or hinted at in this report.

The tools for decision making developed by previous research have already found extensive application in business and government organizations. A number of such applications have been mentioned in this report, but they so pervade organizations, especially at the middle management and professional levels, that people are often unaware of their origins.

Although the research domain of decision making and problem solving is alive and well today, the resources devoted to that research are modest in scale (of the order of tens of millions rather than hundreds of millions of dollars). They are not commensurate with either the identified research opportunities or the human resources available for exploiting them. The prospect of throwing new light on the ancient problem of mind and the prospect of enhancing the powers of mind with new computational tools are attracting substantial numbers of first-rate young scientists. Research progress is not limited either by lack of excellent research problems or by lack of human talent eager to get on with the job.

Gaining a better understanding of how problems can be solved and decisions made is essential to our national goal of increasing productivity. The first industrial revolution showed us how to do most of the world's heavy work with the energy of machines instead of human muscle. The new industrial revolution is showing us how much of the work of human thinking can be done by and in cooperation with intelligent machines. Human minds with computers to aid them are our principal productive resource. Understanding how that resource operates is the main road open to us for becoming a more productive society and a society able to deal with the many complex problems in the world today.

Optimal Human-Computer Task Allocation Via Verifiable Tasks/Methods Matching

Stephen J. Andriole
Department of Information Systems and Systems Engineering
George Mason University

Over the past ten years or so we have designed and developed a variety of decision aids and decision support systems. These projects have largely resulted in the design and development of "prototypes" that have only infrequently been formally evaluated. Many of these systems have been "technology pushed" into development by methodological fads or bureaucratic imperatives that in retrospect have been difficult to always understand.

Of particular concern is the lack of allocation strategies that has contributed to the development of aids and systems that have on occasion badly assigned problem-solving duties between the user and the system. It is possible to identify aids that have failed to even consider optimal task allocation between those intended to use the aids and the aids themselves.

Earlier on in the history of interactive decision aids and support systems, task allocation was of less importance than it is today primarily because our analytical methods were immature. But today—especially via the evolution of knowledge-based methods—task allocation has become a critical design issue. How does a decision aid designer decide where to place problem-solving responsibility in the aiding process? How should methods be selected and matched with tasks? How can genuine methods-based "partnerships" be developed?

One approach to the problem is the development of task taxonomies that characterize problems not just according to their substantive or procedural essence, but according to their amenability to analytical methods and, by implication, their amenability to computer calculation. For example, deductive problems are amenable to many optimization methods, that, in turn, can be almost thoroughly computerized. Inductive problems are much less amenable to complete computer control, while abductive problems may well be completely beyond the state-of-the-art of our analytical methodology and, therefore, computerization. There are other task characteristics that can be used to determine (a) if analytical methods exist to solve them, and (b) if the method(s) can be embedded in computer software requiring very little interaction with the human problem-solver.

Also critical to optimal task allocation is an evolving understanding of what our analytical methods can and cannot do. We have, it seems, over the years developed expectations about our methodology that have been increasingly difficult to justify. Ten years ago decision analysis was touted as the problem-solving answer to the inadequacies of operations research, while today artificial intelligence has about crested in favor of expectations about the power of biological emulation. It is important for us to pause on occasion to identify the kinds of problems to which the methods are well and badly suited. Once these assessments have been made we can begin to match tasks with methods in ways likely to result in optimal task allocation and optimal problem-solving.

Examples of task mis-allocation are all around us. Decision aids and decision support systems equipped with "deep" knowledge bases cannot possibly be expected to infer Soviet intentions toward Poland in the year 2000. At the other extreme is the air traffic control system which insists on using humans to perform tasks for which they are completely unqualified, while computers sit idly by performing trivial tasks.

The presentation will thus address the following questions and issues:

- o what are the characteristics of tasks that will permit assessments about the appropriateness of specific or hybrid analytical methods;
- o what are the inherent strengths and weaknesses of alternative analytical methods classes (like decision analysis, operations research, artificial intelligence, and the like);
- o what approaches to tasks/methods matching are likely to yield optimal combinations and design

blueprints; and

- o what are the design implications of tasks/methods matching, especially regarding the issue of task allocation?

Bibliography

Andriole, S.J. (Ed.). (1985). Microcomputer decision support systems. Wellesley, MA: QED Information Sciences, Inc.

A collection of articles that examine decision support systems (DSS) design approaches, problem-solving methods, and structured evaluation. Some case studies are presented; reviews of popular packages are also included.

Andriole, S.J., and Halpin, S.M. (Eds.). (1986). Information technology for command and control. Special issue of the IEEE transactions on systems, man, and cybernetics. Volume SMC-16, 6, November-December.

A collection of articles arranged in three parts: I, "Information Processing Models for C2"; II, "C2 Information Systems"; and III, "C2 Information Technology Issues and Challenges." Articles deal with methods for capturing expertise, methods for designing and transferring decision aids and decision support systems, and methods for evaluating DSSs. Several articles deal with the man-machine interface (MMI) in interactive systems design and development.

Bonczek, R.H., Holsapple, C.W., and Winston, A.B. (1981). Foundations of Decision Support. New York: Academic Press.

Perspectives on decision support; frameworks for organizational information processing and decision-making; and formal logic approaches to decision support.

Harrison, A.F., and Bramson, R.M. (1982). Styles of thinking. Garden City, New York: Anchor Press/Doubleday.

Brief look at how humans process information for problem-solving with implications for systems design.

Hayes, J.R. (1981). The Complete Problem Solver. Philadelphia, PA: The Franklin Institute Press.

A look at the general processes of problem-solving including memory and knowledge acquisition, decision-making, creativity and invention.

Vick, C.R. and Ramamoorthy, C.V. (1984). Handbook of Software Engineering. New York: Von Nostrand Reinhold Company.

Decision Aids—Who Needs 'Em?

Capt. J.R. Fitzgerald
Commander, Destroyer Squadron 31

J.D. Grossman
Head, Command Systems Div., Navy Personnel R&D Center

It is becoming clear that the victor of the next war will be the side that wins the information war. Warfare is increasingly complicated—covering expanded areas geographically (including the increasing use of space), with more resources to fight the war, more complex interactions between systems and warfighting units and more targets with higher capabilities and greater stealth. Because of the increasing speeds and ranges of enemy weapon systems, future naval warfare commanders must have an up-to-date and complete picture of a very large geographical area. At the same time they must have a current picture of their own resource deployments. The ability to determine and maintain the tactical picture will be one of the technical levers needed to overcome numerical superiority.

A distinction is sometimes drawn between "data" and "information." While data is raw, unprocessed facts or fictions (or pieces of facts or fictions) often made vague by fog of war, information is processed data that enhances the receiver's understanding of his world. Although this is a useful distinction, it can lead to a false sense that information per se is the final frontier. The real objective of winning the information war is to enable a decisionmaker, from the warfare commander on down, to perceive his enemy's intentions, to assess his options accurately and to choose the optimal course of action.

The value of distinguishing between creating information for its own sake and creating it in order to aid in decision making is that the latter objective makes it clear that it is possible to have too much information just as it is possible to have too little. Moreover, the way in which information is presented takes on importance if it affects the quality or timeliness of decision making. Consequently, one of the most potent capabilities needed to win the information war will be the development of decision aids that can handle both situations in which too much information and too little information is available.

This presentation will examine decision aiding from the point of view of the decision maker. To underscore this point and assist the audience in understanding the naval warfare environment for which decision aids are intended, a central theme of this presentation will be a recapitulation of a recent Pacific Fleet Anti-submarine Warfare (ASW) operation conducted by surface ASW forces. During this operation, a number of existing decision aids of varying complexity and sophistication were used by fleet sailors. These aids included the geophysics Fleet Mission Profile Library (FMPL), an Array Heading Rose, Classic Prophet, JOTS 4 and others. The benefits and disadvantages of these aids will be described in the context of the ASW operation. The operational theme is intended to address the question: What does the user need from a decision aid? The objective of the presentation will be to identify research which will culminate in improvements to the utility of decision aids from the user point of view.

A set of user-oriented, top-level requirements will be suggested which need to be validated by researchers and, if validated, will demand more complete definition. According to this scheme, decision aids must be:

1. Synergistic - Permit improved quality and timeliness of decisions.
2. Robust - Resistant to loss of information/Permit graceful degradation in performance.
3. Reliable - Yield consistent results under consistent conditions.
4. Adaptable - Permit variable situations and conditions.
5. Responsive - Capable of self-assessment.
6. Flexible - Permit easy updating.
7. User Friendly - Easy to learn and operate.
8. Testable - Allow thorough testing at all levels of realism and degradation.

In addition to the need to develop decision aids with these characteristics, an understanding of how multiple decision aids with intersecting domains can be integrated to benefit the user is needed.

Finally, a three-dimensional view of the universe of decision aids will be presented in order to identify characteristics of a class of decision aids that can be built, maintained and tested with current technology. One dimension of this model is the extent to which an aid supports a fixed versus a variable decision making environment. Another dimension is the extent to which an aid depends upon organic versus non-organic input. The third dimension is the extent to which the aid supplies a probabilistic output. This view has implications about the extent to which research results have previously been used in the design of decision aids as well as implications for the direction of future research.

Bibliography

Smith, E.E. & Collins, A. (1983). Applied Cognitive Science. Bolt, Beranek and Newman, Inc., Cambridge, MA. ONR Report No. ONR-TR-2.

A discussion of the cognitive aspects of the interaction between people and complex systems. This report presents relatively recent research in the areas of mental models, explanation, models of the user, and engineering.

Andriole, S.J. & Hopple, G.W. (1982, March). They're only human: Decision-makers in command and control. SIGNAL. 36, (7). 45-49.

An argument for designing C2 systems from the decision maker up rather than from the electronics down. This article describes the psychological, psychophysical, and wartime C2 decision making environments and examines the interaction of these environments on the role of the computer. Implications for future research are described.

Greitzer, F.L., Hershman, R.L. & Kaiwi, J. (1985, November). Intelligent interfaces for C2 operability, in: Proceedings of the IEEE International Conference on Systems, Man, & Cybernetics. Tucson, AZ. (pp 782-786).

Argues that the interface to a decision aid should understand the user's goals.

Miller, P.L. (1986). Critiquing: A different approach to expert advice in medicine. Medical Informatics. 11. 29-38.

Discusses the virtues of a decision-aiding approach in which the user of the aid supplies a tentative solution to a problem and the aid critiques the solution.

Zachary, W. (1986, November). Beyond user-friendly: Building decision-aid interfaces for expert end users. Proceedings of the IEEE International Conference on Systems, Man, & Cybernetics. Tucson, AZ. (pp 641-647).

Claims that the designer of a decision aid and the intended user of the decision aid often have different representations of the problem to be solved. Argues that the decision aid's interface must translate between the two representations. Discusses a specific Navy decision aid development as an embodiment of this approach.

Zachary, W. (1986). A cognitively based functional taxonomy of decision support techniques. Human-Computer Interaction. 2. 25-63.

Proposes a generic classification of types of cognitive support needed by decision makers and suggests computer techniques for providing that support.

Decision Aids for Decision Makers: Views of a Decision Analysis Practitioner

D. Warner North
Decision Focus, Incorporated

In the late twentieth-century world of increasing complexity and massive uncertainties, most decision makers need help. Many are aware of this need and are willing to pay handsomely those who can provide them with aids that will help them to deal with their difficult decision situations. Decision analysis has grown over the past several decades as an area of management science specializing in aiding top-level decision makers [1,2]. In some cases the use of quantitative decision aids based on decision analysis and related areas of management science has saved orders of magnitude more than the analysts have charged their decision-maker clients. In other cases, formal decision aids have been perceived as useless or even dangerous by those with decision responsibility.

What can decision research do to help decision makers? There are already a number of powerful concepts and techniques available on which to base decision aids. Some decision aids might be considered as generic: aids to the decision maker and analytical staff to carry out a decision analysis of a broadly defined, generic character. Decision tree software systems and influence diagrams fall into this first category. The second category is decision aids that have been developed to assist on a specific problem, taking advantage of its specific characteristics and often incorporating extensive modeling of complex relationships as well as decision analysis techniques for representing uncertainty. This presentation will illustrate decision aids of both types drawn from the consulting experience of the author and colleagues from SRI International and Decision Focus Incorporated.

Generic Decision Aids

Many decision problems involving one or a sequence of decisions under uncertainty are readily represented as decision trees, and a number of commercial software packages have become available in the last several years for the construction and analysis of decision trees [3]. This software is perhaps most useful as an aid to an experienced decision analyst who is doing an analysis to help a decision maker. For problems involving many stages of cause-effect relationships or an extensive sequence of related events in time, influence diagrams [4,5,6] provide an intuitive means of structuring the uncertainties, as well as a computational method capable of dealing with the equivalent of extremely large decision trees.

Specific Decision Aids

Problems involving complex markets, technical systems, or complex spatial and temporal relationships can often benefit from a decision analysis approach, but extensive modeling specific to the problem may be needed as part of an effective decision aid. One example is the analysis of synthetic fuel commercialization policy done in the mid 1970s for the Ford Administration [7]. (A summary of this analysis is given in the paper by Tani in [2].) This analysis used a model of energy supply/demand relationships to examine the role of synthetic fuels under specific scenarios, and a decision tree to organize the scenarios into an analysis of policy alternatives in the face of uncertainty on foreign oil price, domestic resource availability, and synfuel costs.

A more recent public policy application of decision analysis addressed the need for emissions reduction from power plants to reduce potential damage from acid rain [8]. The character of this decision is whether to act now in the face of considerable uncertainty, or to wait for better information, recognizing that waiting may result in damage that could have been avoided by prompt action. A decision analysis-based software package (ADEPT) has been favorably reviewed [9] and successfully used for a major state-level study by analysts other than those who constructed it [10].

Often decision makers face multiple decisions of the same type. One example is procuring raw materials or fuel supplies in an uncertain market [12]. Another example is the prioritization for testing and treatment among a number of similar environmental problems, such as hazardous waste disposal sites or underground storage tanks.

The latter example has recently been the subject for a decision analysis model that has been applied to tank management within the electric utility industry [13].

Conclusions

Decision analysis has been the basis for successful decision aids, but making the decision aid independent of an experienced analyst is difficult. Generic aids such as decision tree and influence diagram software can make the analyst more efficient, both in doing analysis and in communication with the decision maker and his or her advisors. The acid rain and underground tanks examples suggest that there is a high potential for computer-implemented decision aids that provide concepts and computational techniques from decision analysis tailored to a specific class of decision problems. Such tools may facilitate the use of a sophisticated decision analysis approach without the extensive involvement of an experienced analyst, which has until recently characterized most successful decision analyst applications.

Bibliography

- [1] Keeney, Ralph L. (1982). Decision Analysis: An Overview. Operations Research. 30 (5). 803-838.
- [2] Howard, Ronald A. and Matheson, James E. (Eds). (1983). The Principles and Applications of Decision Analysis. Menlo Park, CA: Strategic Decision Group.
- [3] Henrion, Max. (1985). Software for Decision Analysis: A Review of Riskcalc, Arborist, and Supertree. OR/MS Today. 12 (5). 24-29.
- [4] Miller, A.C., III et al. (1976, May). Development of Automated Aids for Decision Analysis. Report prepared by SRI International for the Defense Advance Research Projects Agency.
- [5] Howard, Ronald A. and Matheson, James E. (1983). Influence Diagrams. The principles and applications of decision analysis. (pp 719-782). Menlo Park, CA: Strategic Decisions Group.
- [6] Schacter, Ross. Evaluating Influence Diagrams. Operations research. in press.
- [7] Synfuels Interagency Task Force. (1975). Recommendations for a synthetic fuels commercialization policy. Washington, D.C.: U.S. Government Printing Office.
- [8] North, D. Warner, Balson, William & Colville, Glenn. (1986). Representing uncertainty explicitly in environmental risk management: An example from acid deposition policy analysis. Proceedings of an APCA International Specialty Conference: Environmental Risk Management: Is Analysis Useful? Pittsburgh, PA: Air Pollution Control Association.
- [9] Acid Deposition Decision Tree (ADEPT) Software Review. (1985, December). Microcomputer software reviews. Journal of the Air Pollution Control Association. 35 (12). 1312-1313. December 1985.
- [10] Bratcher, R. (1986, October). The application of risk management methods to the acid rain debate in Wisconsin. Presented at the Harvard Symposium on Environmental Risk Management: Research Needs and Opportunities, John F. Kennedy School of Government, Harvard University, Cambridge, MA.
- [11] Tech transfer news: Washington state uses ADEPT. (1987). EPRI Journal 12 (1). 38-39.
- [12] Fancher, Richard B., Wilson, James, and Mueller, Howard A. (1985). Fuel contracting under uncertainty. Paper presented at the Summer Meeting of the IEEE Power Engineering Society.
- [13] Smith, A., Cohan, D., et al. (1987). The underground tank risk management model: A guide to the TANKS decision support software. Report prepared for the Electric Power Research Institute by Decision Focus Incorporated.

Interpretation-Based Decision Aids

Karl Weick
University of Texas - Austin

Good decisions are more likely when meanings are clear and environments are stable. To make a good decision, people need a clear sense of what problem needs to be solved and which options have which reliable effects. These issues can be treated as problems in sensemaking and meaning. From this perspective, aids to decision making often take the form of devices that aid interpretation (Daft & Weick, 1984).

Weick (1985) reviewed five devices for sensemaking used by organizational members—enactment, triangulation, interaction, deliberation, and abstraction—and argued that conventional decision aids often make it more difficult for people to use these five devices when they try to understand what needs to be decided. It will be suggested that if these five sources of sensemaking are more systematically used in the context of decision making, decisions can be improved. Decision aids involving action rationality, future perfect thinking, minority influence, stress reduction, and narrative rationality will be used to illustrate the argument.

Before we explore these five sensemaking devices, a general point should be made. Interpretation aids often work as much because of the processes of reflection and examination they trigger as because of the specific manipulations of content they perform. For example, a useful exercise to build a stable picture from which problems and decisions can be inferred is to construct a cause map (e.g., Hall, 1984). Cause maps are a means to reduce uncertainty produced by disagreement about cause-effect relations (e.g., Thompson, 1967). While the cause map generated during group discussions often improves decision making because it organizes variables by their causal implications, the discussions themselves induce a critical mindset so that people are then able to see more clearly what needs to be decided apart from what is codified in a map.

Many decision aids may have their greatest impact through these second-order effects. This suggests that one decision aid is better than no decision aid, but that one decision aid may be just as good as any other aid, because the improvement comes from the discussion that the aid triggers, not from the way the aid itself structures information.

If we turn now to the five bases of sensemaking and to the question of what decision aids they suggest, we start with enactment, the process by which people learn about events by prodding them. A growing number of researchers (e.g., Starbuck, 1985) argue that direct action is a form of deliberation and, therefore, that action and thought should be interspersed with one another more frequently than they now are and that action should be introduced in any episode of deliberation sooner than it usually is. Action generates environmental stability, current information, and something tangible which can edit potential choices. Thus, analysis of decision aids such as a tree diagram should be abbreviated and interrupted systematically by physical action which introduces new data to the problem and allows people to "act first, think later."

The second process of sensemaking, triangulation, involves the systematic use of alternative data sources and is illustrated by future perfect thinking. Decision aids typically focus on contemporary problems. While they occasionally stretch away from the present into the future and the past, these temporal extensions are usually modest. If we relax the present as a constraint, then we find additional ways in which decision performance can be aided. Boland (1984), for example, gathered a group of film lending executives in 1980, provided them with accounting reports prepared for 1982 to 1985, asked them to imagine it was July 21, 1985, and then to discuss what the film service had become and why. This exercise in future perfect thinking was an attempt to explore the proposition that it is easier to make sense of events when they are placed in the past, even if the events have not yet occurred. Boland reported that a major outcome of the experiment was that, in trying to understand what had been done in an imaginary future, participants discovered that they had an inadequate understanding of an actual past. The experiment uncovered disagreements about the nature and meaning of past events that people did not realize had impeded their current decision making.

The point of the Boland work, and the more general concept of future perfect thinking (Weick, 1979), is that sensemaking can be extended beyond the present. As a result, present decisions can be made meaningful in a larger context than they usually are and more of the past and future can be brought to bear to inform them.

The third sensemaking device, social interaction, is invoked when people learn about events by comparing what they see with what someone else sees and then negotiating some acceptable version of what really happened. Suggestions for decision aids that are sensitive to this mechanism emerge from Nemeth's (1986) recent research on minority influence. She suggests that minority viewpoints are important because they stimulate divergent attention and thought. People exposed to opposed minority views "are stimulated to attend to more aspects of the situation, they think in more divergent ways, and they are more likely to detect novel solutions or come to new decisions." Decisions, as a result, are expected to be better or more accurate. Minority positions affect thought processes, which means that interpretation-based decision aids should be designed to introduce credible minority viewpoints.

The fourth sensemaking device, deliberation, is invoked when people learn about events through slow and careful reasoning. While this device is the core activity associated with most decision aids, its functioning is sensitive to variations in threat, uncertainty, and unpredictability. Decision aids need to address deliberation under pressure since pressure has such marked effects on information processing. A good example of an interpretation-based decision aid derives from the well-established finding that when people think under pressure they revert to over-learned response tendencies, experience perceptual narrowing, and centralize decision making (Staw, Sandelands, & Dutton, 1981). These three responses to strain increase the probability that people will make outdated choices for environments that no longer exist using criteria that are irrelevant. Decision aids could neutralize these outcomes if they lowered the stress itself (reduce importance of situation, reduce demands, raise ability to cope with demands), simplify the task being performed, compensate for perceptual narrowing, or increase the person's access to more recently learned complex responses.

The fifth sensemaking device, abstraction, occurs when people understand an event by building a context around it, an activity which often requires that they move to a higher level of abstraction. Such a move toward embedding a poorly understood event in a larger context is suggested by recent work on narrative rationality (Robinson & Hawpe, 1986). These studies suggest that if people could study a portfolio of organizational stories before they made decisions, better decisions would be made. Stories are mnemonic devices that help people recall larger, more complex sets of decision premises, assumptions, and experiences before they actually make a decision. This recall should improve decision making directly, because more information is brought to bear on the decision, and indirectly because priming through stories creates the mindset of narrative rationality. Narrative rationality introduces considerations of coherence, plausibility, fit with past experience, temporal development, and resolution, criteria which are often less salient when the mindset of traditional argumentative rationality is evoked.

Bibliography

- Boland, J.R. (1984). Sense-making of accounting data as a technique of organizational diagnosis. Management Science. 7 868-882.
- Daft, R.L., & Weick, K.E. (1984). Toward a model of organizations as interpretation systems. Academy of Management Review. 9, 284-295.
- Hall, R.I. (1984). The natural logic of management policy making: Its implications for the survival of an organization. Management Science. 30, 905-927.
- Nemeth, C.J. (1986). Differential contributions of majority and minority influence. Psychological Review. 93 23-32.
- Robinson, J.A., & Hawpe, L. (1986). Narrative thinking as a heuristic process. In: T. Sarbin (Ed.), Narrative Psychology. (pp 111-125). New York: Praeger.
- Starbuck, W.H. (1985). Acting First and Thinking Later. In J. Pennings (Ed.), Organizational Strategy and Change (pp 336-372). San Francisco: Jossey-Bass.
- Staw, B.M., Sandelands, L.E., & Dutton, J.E. (1981). Threat rigidity effects in organizational behavior: A multi-level analysis. Administrative Science Quarterly. 26 510-524.
- Thompson, J.D. (1967). Organizations in Action. New York: McGraw-Hill.
- Weick, K.E. (1979). The Social Psychology of Organizing. 2nd Edition. Reading, MA: Addison-Wesley.
- Weick, K.E. (1985). Cosmos vs. Chaos: Sense and Nonsense in Electronic Contexts. Organizational Dynamics. 14 50-64.